

DB2 Universal Database Server for OS/390 and z/OS



What's New?

Version 7

DB2 Universal Database Server for OS/390 and z/OS



What's New?

Version 7

Note

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 53.

Second Edition (May 2001)

This edition applies to Version 7 of IBM DATABASE 2 Universal Database Server for OS/390 and z/OS (DB2 for OS/390 and z/OS), 5675-DB2, and to any subsequent releases until otherwise indicated in new editions. Make sure you are using the correct edition for the level of the product.

This and other books in the DB2 for OS/390 and z/OS library are periodically updated with technical changes. These updates are made available to licensees of the product on CD-ROM and on the Web (currently at www.ibm.com/software/data/db2/os390/library.html). Check these resources to ensure that you are using the most current information.

This softcopy version is based on the printed edition of the book.

© **Copyright International Business Machines Corporation 2001. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this book	v
Chapter 1. The soul of e-business	1
Using DB2 for OS/390 and z/OS as your Web server	1
Application servers on the Web	3
Database access with DB2 Connect	3
Transaction manager support for multisite updates	4
High-speed text search for your Web applications	5
DB2 for OS/390 and z/OS Version 7 for your e-business and business	
intelligence applications	5
Warehouse integration	5
Improved performance, availability, and scalability	5
More power for your enterprise applications.	7
Chapter 2. What's new in DB2 for OS/390 and z/OS Version 7?	9
More power for managing data in DB2 for OS/390 and z/OS	9
Comprehensive statistics history	9
Utility lists with pattern matching and dynamic allocation	9
More DBADM authority	10
Enhanced management of constraints	10
Deferred definitions of data sets	10
DB2 for your e-business — reliable, scalable, and available	10
A new UNLOAD utility	10
A new COPYTOCOPY utility	11
Parallel LOAD with multiple inputs	11
Improved optimization and parallelism	12
Faster online REORG	13
More concurrency with online LOAD RESUME	13
Changing system parameters without stopping DB2	13
Improved availability of some user objects	13
Improved availability with CHECKPAGE for COPY	14
Log suspend and resume for external copy	14
Allow retry of log-read request	14
DDF suspend and resume.	14
Index access for small tables.	14
Data sharing enhancements	14
Outstanding performance when running DB2 on Enterprise Storage Server.	16
Easier integration of enterprise applications	16
DB2 XML Extender for OS/390 and z/OS for your e-business.	16
Unions in views.	18
More flexibility with SQL	18
Global transaction support for distributed applications.	20
Better monitoring of the distributed environment.	20
Support for loading output from a SELECT Statement	20
Support for Java stored procedures	20
Stored procedures that are written entirely in SQL	21
ODBC 3.0 support	21
Savepoints	21
Support for defining identity columns	21
Ability to define declared temporary tables.	21
SQL support for REXX	21
New Precompiler Services.	22
Unicode support for international data stores and e-business	22

Enhanced DB2 Extender for text objects	22
Enhanced business partnerships with independent software developers	23
Improvements in connectivity.	23
Support for COMMIT and ROLLBACK in stored procedures	23
Support for Windows Kerberos security	23
Reporting server elapsed time at the workstation	24
New packaging for DB2 utilities	24
Support for Migration to Version 7 of DB2 for OS/390 and z/OS from Version 5 or Version 6	25
Information to help you use DB2	25
New books about DB2 for OS/390 and z/OS	25
DB2 for OS/390 and z/OS on the Web	26
Helpful Web sites	26
DB2 family	26
Other references	27
Chapter 3. Managing your Enterprise	29
Integrating and analyzing your business information using the DB2 Warehouse Manager feature	29
DB2 Warehouse Manager	29
Reporting and governing your enterprise using QMF	31
Managing your enterprise with the DB2 Management Clients Package	34
Managing DB2 for OS/390 and z/OS from your workstation with Control Center	34
Building DB2 stored procedures from your workstation	37
Installing DB2 from a workstation	38
Using workstation views of DB2 Explain output	39
Estimating DB2 performance with DB2 Estimator	41
Web-ready applications.	42
Text search capability for the Internet with Net Search Extender	42
Net.Data for secure Web applications	44
Data Management Tools for your database	44
Glossary	49
Notices	53
Trademarks	54
Index	57

About this book

This book provides an executive overview of Version 7 of IBM® DB2 Universal Database™ for OS/390®.

This book refers to the DB2® family as DB2 Universal Database. This book refers to IBM DB2 Universal Database Server for OS/390 as *DB2 for OS/390 and z/OS*. In cases where the context makes the meaning clear, this book uses the term *DB2* in place of the longer references.


The term *MVS* represents the MVS/Enterprise Systems Architecture (MVS/ESA™) element of OS/390. The term *CICS*® represents CICS/ESA® and CICS Transaction Server for OS/390. *IMS*™ represents IMS/ESA®. *C* and *C language* are terms that represent the C programming language.

Although the focus of this book is on Version 7 of DB2 for OS/390 and z/OS, IBM has also made many of the functions and capabilities available to existing Version 6 users. This book identifies these functions and capabilities with the following pair of graphic symbols.



V6

Information about some function or capability that is available in Version 7 and Version 6.



Users of DB2 for OS/390 Version 5, Version 6, or Version 7, or people who want to know about DB2, can subscribe to iSource, an IBM service that sends notification of announcements about DB2 and other IBM software. For more information about how to subscribe to this service, see the following Web site:

ibm.com/isource

Version 5 and Version 6 users can learn how to obtain some capabilities now.

ibm.com/software/db2os390/v5apar.html

ibm.com/software/db2os390/v6apar.html

Chapter 1. The soul of e-business

Are you ready for e-business? A different world is out there. It's a world where eight seconds is too long for a customer to wait, and one minute of downtime costs thousands of dollars. Web users speak the languages of the world. Business models and winning applications change every day.

In some ways, e-business is very simple. Customers want the best of everything, and they want it now. If yours is not the best, then the competition is just a mouse click away. For e-business applications, the demands are similar to your traditional applications: fast, scalable, and available, with flexible access to usable information.

The traditional transaction processing strengths of the DB2 family, along with exciting new functions, deliver what you need for e-business. Whether your customers use Windows 2000, UNIX®, Linux, OS/390, or z/OS operating systems, you can deliver data with DB2. Hardware choices are just as versatile, ranging from the new wireless handheld devices like Palm VII to the most powerful workstations and servers, such as the new zSeries 900 e-business enterprise server with 64-bit architecture. DB2 software makes traditional system barriers irrelevant.

Using DB2 for OS/390 and z/OS as your Web server

Many businesses choose to use DB2 for OS/390 and z/OS as their database server for the Web. The three-tier architecture is ideal for many e-businesses. In a three-tier architecture:

- The client is on the first tier (as it is in a two-tier architecture). The client handles user interactions; it does not contain any business logic. A second, or middle, tier supplements the client.
- One or more application servers reside on the middle tier. The application server handles the portion of the business logic that does not require the functionality that the database server provides.
- On the third tier is the database server, DB2 for OS/390 and z/OS, that provides data for the application. All Web-based database applications rely on a database server. For performance reasons, the database server typically uses stored procedures to handle some of the business logic.

The three primary components of Web-based database applications are the Web browser (or client), the Web application server, and the database server. The client handles the presentation logic and, in some cases, validates user-provided input. Web applications sometimes integrate Java applets or JavaScript into the client-side logic to improve the presentation layer. Hardware and software components of the second and third tiers share responsibility for the availability, scalability, and performance characteristics of the Web site.

In many cases, the client and server for a Web application are on different operating systems. The client, for example, can be on a workstation-based operating system, such as Windows 2000. The server for the application can also be on a workstation-based operating system or on an enterprise server operating system, such as OS/390 and z/OS.

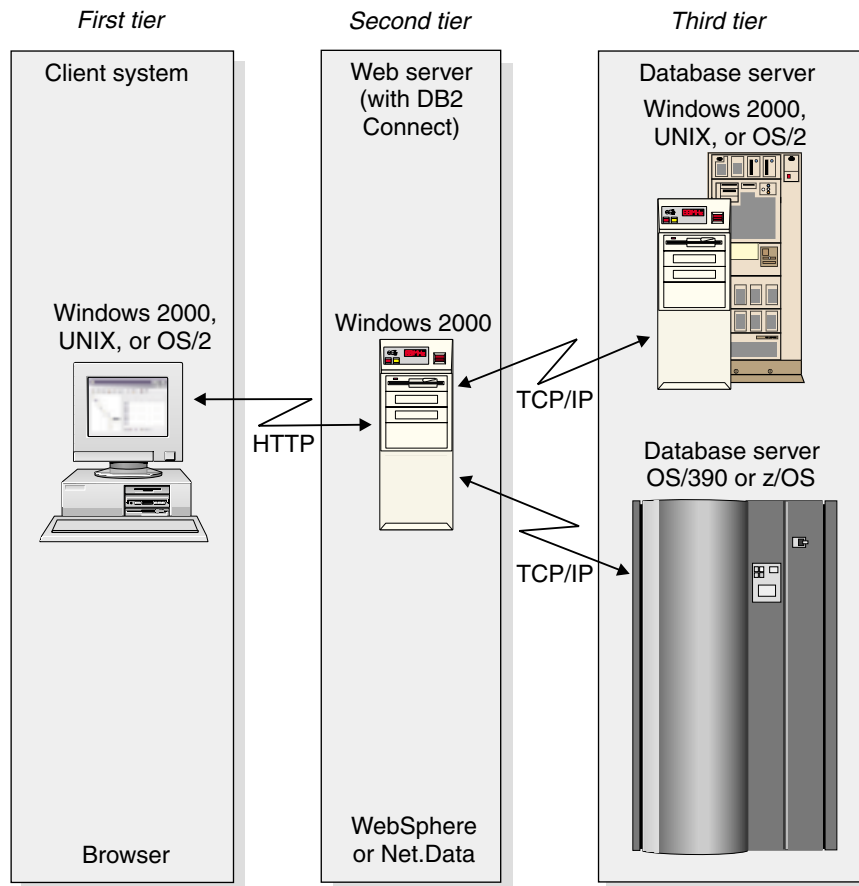


Figure 1. Three-tier connectivity with a workstation-based client and Web server, and with different database servers

Figure 1 illustrates the use of a three-tier architecture. In this example, the Web server on the second tier is a Windows 2000 machine. WebSphere or Net.Data[®] can use DB2 Connect to access DB2 for OS/390 and z/OS. You will read more about WebSphere, Net.Data, and DB2 Connect later in this section.

DB2 for OS/390 and z/OS offers a robust solution for Web applications. Specifically, using DB2 for OS/390 and z/OS as a database server for a Web application provides the following advantages:

- **Tremendous scalability.** The volume of transactions on any Web application varies. Transaction loads increase, or spike, at different times of the day, on different days of the month, or at different times of the year. Transaction loads also tend to increase over time, as more users move to the Web. In a Parallel Sysplex environment, DB2 for OS/390 and z/OS can handle the full range of transaction loads with little or no impact on performance. Any individual user is generally unaware of how busy the system is at a given point in time.
- **High degree of availability.** When DB2 for OS/390 and z/OS runs in a Parallel Sysplex environment, the availability of data and applications is very high. If one DB2 subsystem is unavailable because of maintenance, for example, other DB2 subsystems in the Sysplex take over the workload. Users are unaware that part of the system is unavailable because they have access to the data and applications that they need.

- **Ability to manage a mixed workload.** DB2 for OS/390 and z/OS effectively and efficiently manages priorities of a mixed workload as a result of its tight integration with OS/390 Workload Manager (WLM).
- **Protection of data integrity.** Users of DB2 for OS/390 and z/OS can benefit from the product's well-known strength in the areas of security and reliability.

Application servers on the Web

IBM provides two application servers, WebSphere and Net.Data, that help companies "Web-enable" their data and business logic. WebSphere and Net.Data products run on popular operating systems, including AIX®, Linux, OS/390, OS/400®, Windows 2000, Windows NT®, and Solaris Operating Environment.

WebSphere

The WebSphere family of products offers users:

- A Java server run-time environment that is based on industry standards
- Web-site development tools
- Commerce software
- Management software

By using these tools, companies can build, deploy, and manage portable e-business applications. Information about WebSphere is available at:

ibm.com/software/webservers

Net.Data

The Net.Data family of products provides a robust environment for both application development and execution. Companies can use Net.Data products to create high-performance Web applications by using a variety of programming languages, such as Java, REXX, Perl, and C++. IBM delivers more power with Net.Data for Version 7 of DB2 for OS/390 and z/OS. See "Net.Data for secure Web applications" on page 44 for more information about Net.Data.

Database access with DB2 Connect

DB2 Connect gives applications fast and easy access to existing databases on IBM enterprise servers. It has a highly scalable communication infrastructure for connecting Web, Windows®, UNIX (including Linux), OS/2® and mobile applications to S/390 and AS/400® data. For detailed information about DB2 Connect, see *DB2 Connect User's Guide*.

DB2 Connect provides extensive application programming tools for developing client-server and Web applications using industry standard APIs such as ODBC, JDBC™, and SQLJ.

DB2 Connect Version 7 includes the following features that enhance its database access capabilities:

- Improved scalability and performance with connection concentration
Connection concentration, also referred to as transaction pooling, benefits applications that acquire a connection and retain the connection for long periods of time. Connection concentrator technology concentrates the workload from all applications in fewer S/390 host or AS/400 database server connections. The connection concentrator provides:
 - Automatic thread reuse following commit and rollback points
 - Delayed thread reuse when an open WITH HOLD cursor is detected
 - Intact SQL special registers across transaction boundaries

- Improved load-balance and fail-over support for DB2 for OS/390 and z/OS servers in Parallel Sysplex® environments
- Distributed request capability that lets you work with all DB2 family servers in the same SQL statement, and access Oracle databases by using the DB2 Relational Connect option.
- Support for new function in DB2 for OS/390 and z/OS Version 7 that includes:
 - Scrollable cursors, described in “Scrollable cursors” on page 18.
 - FETCH FIRST *n* ROWS ONLY, described in “FETCH FIRST *n* ROWS SQL clause and fast implicit close” on page 19.
 - Single sign-on capability using Kerberos and user ID and password encryption, described in “Support for Windows Kerberos security” on page 23.
- Improved transactional capabilities with support for tightly-coupled X/Open XA transactions, such as those required by BEA Tuxedo, and support for Microsoft® SNA Server in two-phase commit applications

Transaction manager support for multisite updates

In a *multisite update*, data is updated in more than one database within a single transaction. A distributed transaction typically involves multiple recoverable resources, such as DB2 tables, MQSeries® messages, and IMS databases. With *global transaction support*, new in Version 7, independent DB2 agents can participate in a single global transaction, share locks, and access the same data within a single transaction, as described in “Global transaction support for distributed applications” on page 20.

A transaction manager coordinates commit operations across a distributed transaction. DB2 supports a wide range of transaction manager products. Application environments that use DB2 Connect® to access DB2 remotely can use the following transaction manager products:

- Transaction manager products that support standard XA protocols
- IBM WebSphere™
- BEA Tuxedo
- BEA WebLogic
- Microsoft Transaction Server (MTS)
- IBM TxSeries (CICS and Encina)
- Java applications that support Java Transaction API (JTA) and Enterprise JavaBeans (EJBs)
- Message and Queuing Series (MQSeries)

The XA protocol is a bidirectional interface between a transaction manager and resource managers that provides coordinated updates across a distributed transaction. In this environment, DB2 Connect, DDF, and local attachment facilities (RRSAF, CICS, IMS™) now recognize database connections that are part of a global transaction. Locks are shared across the branches of a distributed global transaction.

Example: An application server initiates a global transaction and uses Enterprise JavaBeans™ to invoke applications at other application servers. These other application servers then invoke additional applications at other servers as part of the global transaction.

When the application runs, the middleware, such as DB2 Connect, identifies the various branches of the global transaction so that DB2 (as a resource manager)

understands the relationships between the database connections. DB2 automatically lets the branches of a global transaction within a given DB2 subsystem share locks so that one branch does not lock out other branches of the same global transaction.

High-speed text search for your Web applications

A new feature of DB2 for OS/390 and z/OS, Net Search Extender, delivers the latest in high-speed Internet text search capability. This feature addresses the needs of e-commerce and Web self-service applications. The extender handles the heavy text search demands of large text-intensive Web sites by searching text data in DB2. Net Search Extender rapidly searches data without locking database tables and delivers excellent query performance and scalability using in-memory database technology. For more information, see "Text search capability for the Internet with Net Search Extender" on page 42

DB2 for OS/390 and z/OS Version 7 for your e-business and business intelligence applications

With Version 7, DB2 for OS/390 and z/OS delivers improved performance, availability, and scalability for your e-business and data warehouse applications. Using the powerful environment provided by System/390®, OS/390, and z/OS, you can leverage your existing applications while developing and expanding your electronic commerce for the future.

This section provides a quick overview of the new function. For more detailed information on each of these enhancements, see "Chapter 2. What's new in DB2 for OS/390 and z/OS Version 7?" on page 9 and "Chapter 3. Managing your Enterprise" on page 29.

Warehouse integration

The DB2 family is expanding the capability of the DB2 server by integrating warehouse management. The new DB2 Warehouse Manager makes it easy to design and deploy a data warehouse on your S/390. You can extract operational data from your DB2 for OS/390 and z/OS and import it into an S/390 warehouse without transferring your data to an intermediate platform. You can prototype applications more quickly, query and analyze data, and help your users access data and understand information. The new DB2 Warehouse Manager feature gives you a full set of tools for building and using a data warehouse based on DB2 for OS/390 and z/OS.

Improved performance, availability, and scalability

The DB2 for OS/390 and z/OS Version 7 delivers improved performance, availability, and scalability.

DB2 utilities

DB2 has included a basic set of core utilities since Version 1. These utilities initially provided a basic level of services that let customers manage data. Some customers, however, preferred to obtain such functions from independent software vendors that develop utilities and tools that provided additional performance, function, and features beyond those contained in DB2 utilities. In response to clear customer demand, IBM has invested in improving DB2 utilities in recent releases of DB2 for OS/390 and z/OS.

IBM now offers DB2 utilities as products in the new family of Data Management Tools for DB2. These new DB2 utility products deliver improved and powerful tools for your database operation and also for diagnosis and recover. For more information, see “New packaging for DB2 utilities” on page 24.

New UNLOAD utility

A new utility, UNLOAD, provides faster data unloading than was available with the DSNTIAUL program. In addition to the unload functions of REORG UNLOAD EXTERNAL, you can also unload data from an image copy, do row sampling, specify field selection and ordering, and use parallelism when unloading multiple partitions.

LOAD utility enhancements

The LOAD utility has three major improvements:

- You can use SQL SELECT to produce input directly into LOAD. This enhancement delivers a DB2 family cross loader function. Use any DRDA-compliant server as a data input source for populating your tables in DB2 for OS/390 and z/OS. Your input can even come from non-DB2 data sources by using IBM DataJoiner for access to data from sources as diverse as Oracle and Sybase, as well as the entire DB2 family of database servers. See “Support for loading output from a SELECT Statement” on page 20.
- In earlier releases, many customers would run individual LOAD utilities at the partition level in order to take advantage of multiple processors. Now this same function can be performed with a single LOAD job with multiple partitions loaded in parallel and with added efficiency building secondary indexes. See “Parallel LOAD with multiple inputs” on page 11.
- You can specify SHRLEVEL CHANGE on the LOAD utility for improved data availability. You enable LOAD to run concurrently with queries and applications.

REORG utility enhancements

Online REORG makes your data more available. Enhancements to this utility eliminate the time required to rename datasets and shorten the elapsed time. When hundreds of table space and index objects are involved in a reorganization, the time savings can become quite significant.

New COPYTOCOPY utility

The COPYTOCOPY utility makes up to three additional backup copies asynchronously from an existing copy and registers the copies in the DB2 catalog for recovery purposes.

Other availability enhancements

In addition to more powerful utilities, DB2 for OS/390 and z/OS Version 7 improves availability with the following enhancements:

Change subsystem parameters without stopping DB2

With Version 7 you can change the value of many of subsystem parameters without stopping DB2.

Restart light for data sharing

Data sharing customers can benefit from a new restart option. Restart light lets you restart DB2 with a smaller storage footprint so that you can quickly recover retained locks following an abnormal termination.

More power for your enterprise applications

Version 7 gives you more flexibility for your enterprise applications and makes it easier to call applications.

DB2 XML Extender for OS/390 and z/OS

The family adds DB2 XML Extender for OS/390 and z/OS support for the XML data type. This extender lets you store and retrieve an XML object in either an XML column for the entire document, or in several columns that contain the fields from the document structure.

Improved scrollable cursors

In previous releases of DB2, cursors scroll sequentially through a result table or a result set. With Version 7, scrollable cursors can move forward, backward, or randomly through a result table or a result set. As a result, your application no longer needs to cache the resultant data or to reinvoke the query to reposition the cursor within the result data.

Support for Unicode encoded data

In the increasingly global world of business and e-commerce, there is a growing need for data arising from geographically disparate users to be stored in a central server. Previous releases of DB2 have offered support for numerous code sets of data in either ASCII or EBCDIC format. However, there was a limitation of only one code set per table. New in Version 7 is support for Unicode encoded data. This new code set is an encoding scheme that can represent the characters (code points) of many different geographies and languages.

Stored procedures enhancement

Stored procedures introduced in Version 5 have increased program flexibility and portability among relational databases. Now, with Version 7, DB2 for OS/390 accepts COMMIT and ROLLBACK statements that are issued within a stored procedure. This enhancement is especially useful for applications in which the stored procedure has been invoked from a remote client.

Support for Precompiler Services

Now, you can take advantage of support for Precompiler Services to perform the tasks currently executed by the DB2 precompiler. By using this option, you can eliminate the DB2 precompile step in program preparation and take advantage of language capabilities that had been restricted by the precompiler. Using the host language compiler enhances DB2 family compatibility, and makes it easier to import applications from other database management systems and from other operating environments.

SQL enhancements

Greater flexibility and family compatibility come from several SQL enhancements. Now, you can use a subselect to determine the values that were used in the SET clause of an UPDATE statement. Also, you can have a self-referencing subselect. The search condition in the WHERE clause can include a subquery in which the object for the subquery and the UPDATE or DELETE is the same. You can define a view with the operators UNION or UNION ALL. You can also use a row expression to compare with a subselect expression.

Chapter 2. What's new in DB2 for OS/390 and z/OS Version 7?

IBM announces enhancements to DB2 for OS/390 and z/OS that support your needs on System/390. DB2 for OS/390 and z/OS Version 7 delivers improved performance, availability, and scalability for data warehouse applications and e-business.

The DB2 family expands the capability of the DB2 server by integrating warehouse management. You can integrate your e-business and business intelligence applications by using the powerful, highly available environment provided by System/390, OS/390 and z/OS. You can leverage your existing applications while developing and expanding your electronic commerce for the future.

More power for managing data in DB2 for OS/390 and z/OS

Several new facilities in DB2 for OS/390 and z/OS help you track changes, submit utility changes, manage authorizations, and share data across your enterprise.

Comprehensive statistics history

As the volume and diversity of your business activities grow, you need to change the physical design of DB2 objects. Version 7 of DB2 for OS/390 and z/OS collects statistics history to track your changes. With historical statistics available, DB2 can predict the future space requirements for table spaces and indexes more accurately and run utilities to improve performance. DB2 Visual Explain uses statistics history as a base against which to compare new variations that you enter so that you can improve your access paths.

DB2 stores statistics in catalog history tables. To maintain optimum performance of processes that access these tables, use the MODIFY STATISTICS utility. The utility can delete records that were written to the catalog history tables based on a specific data or that are those records that are a specific age.

Utility lists with pattern matching and dynamic allocation

With Version 7, database administrators can submit utilities jobs more quickly and easily. Now you can:

- Dynamically create object lists from a pattern-matching expression
- Dynamically allocate the data sets required to process those objects

Using a LISTDEF utility control statement, you can standardize object lists and the utility control statements that refer to them. Standardization reduces the need to customize and change utility job streams over time. See Figure 2 for an example of how you can define a list that you can pass to one or more utilities.

```
...  
LISTDEF payroll INCLUDE TABLESPACE PAYROLL.*  
              INCLUDE INDEXSPACE PAYROLL.*IX  
              EXCLUDE TABLESPACE PAYROLL.TEMP*  
              EXCLUDE INDEXSPACE PAYROLL.TMPIX*  
COPY LIST PAYROLL ...
```

Figure 2. Example of a LISTDEF utility control statement that specifies a list of objects.

The TEMPLATE utility control statement simplifies job control language (JCL) by eliminating most data definition (DD) cards. Now you can provide data set templates, and DB2 dynamically allocates the data sets that are required based on your allocation information.

With these utility control statements, database administrators require less time to maintain utilities jobs, and database administrators who are new to DB2 will learn to perform these tasks more quickly. These factors can help reduce the cost of operations.

More DBADM authority

Database administrators now have more flexibility to administer privileges. A person with DBADM authority can create views for others using Version 7 of DB2 for OS/390 and z/OS. This change lessens the need for granting SYSADM authority.

Enhanced management of constraints

You can specify a constraint name at the time you create primary or unique keys. DB2 introduces the restriction of dropping an index that is required to enforce a constraint.

V6

Deferred definitions of data sets

You can defer defining data sets when you have many tables that do not have data in them. This enhancement enables faster installation of applications and makes data set management easier for some CRM and ERM products such as SAP R/3 and PeopleSoft products.

DB2 for your e-business — reliable, scalable, and available

Many businesses already use DB2 for their e-business. They depend on the reliability, scalability, and availability of DB2 to serve their worldwide customers around the clock, seven days a week. Version 7 offers many improvements for electronic commerce. If you have not yet deployed your business on the Web, now is the ideal time to start. A number of enhancements in Version 7 of DB2 for OS/390 and z/OS help you gain competitive advantage.

A new UNLOAD utility

You can unload data from a table space or an image copy data set by using the new UNLOAD utility. In most cases, the UNLOAD utility is faster than the DSNTIAUL sample program, especially when you activate partition parallelism. UNLOAD is easier to use than REORG UNLOAD EXTERNAL. As shown in Figure 3 on page 11, using UNLOAD you can perform the following tasks:

- Unload data from an image copy data set.
- Unload data from multiple partitions in parallel.
- Select data by using a syntax similar to the SQL SELECT statement.
- Sample rows by table.
- Use field selection, ordering, and formatting options.
- Specify SHRLEVEL CHANGE or REFERENCE.




Figure 3. Unloading data using the new DB2 UNLOAD utility

A new COPYTOCOPY utility

While the COPY, LOAD, and REORG utilities can make two local and two remote site backup copies of data, you might want to make a single copy and then clone that copy at a more convenient time. The COPYTOCOPY utility makes up to three additional backup copies asynchronously from an existing copy and registers the copies in the DB2 catalog for recovery purposes. You can use object wildcarding and dynamic allocation capabilities with the COPYTOCOPY utility.

Parallel LOAD with multiple inputs

Using Version 7, you can easily load large amounts of data into partitioned table spaces for use in data warehouse or business intelligence applications. Parallel load with multiple inputs runs in a single step, rather than in different jobs.

The LOAD utility loads each partition from a separate data set so that one job can load multiple partitions in parallel. Parallel loading reduces the elapsed time for

loading the data as compared to loading the same data with a single job in earlier releases. Using load parallelism is much easier than creating multiple LOAD jobs for individual parts.

Figure 4 shows a parallel load of four partitions, with the SORTKEYS keyword enabling a parallel index build of three indexes. Each load task takes input from a sequential data set and loads the data into a corresponding partition. The utility then extracts index keys and passes them in parallel to the sort task that is responsible for sorting the keys for that index. If there is too much data to perform the sort in memory, the sort product writes the keys to the sort work data sets. The sort tasks pass the sorted keys to their corresponding build task, each of which builds one index. If the utility encounters errors during the load, DB2 writes error and error mapping information to the error and map data sets.




Figure 4. Parallel LOAD with multiple inputs and parallel index build

Improved optimization and parallelism

Several enhancements deliver more efficient processing of SQL queries and better performance for complex queries.

More efficient processing for SQL queries

The way DB2 for OS/390 and z/OS Version 7 processes queries can improve performance of your SQL queries. Among the improvements are the following items:

- Fewer sort operations for queries that have an ORDER BY clause and WHERE clauses with predicates of the form COL=*constant*
- IN-list index access parallelism, which can improve performance for queries involving IN-list index access
- The ability to use indexes with a correlated subquery so that ERM applications, such as PeopleSoft applications, will perform and scale much better

V6

Better optimization for complex queries

DB2 for OS/390 and z/OS can use a new star join method. The method provides better optimization and execution performance for queries that join tables together in a star schema. A *star schema* consists of a fact table and a number of dimensions, each of which consists of one or more tables. In addition to improving execution performance, this enhancement enables DB2 to handle more complex star schemas and more tables in a join.

Faster online REORG

Online REORG enhancements improve the availability of data in two ways:

- Online REORG no longer renames data sets, which greatly reduces the time that data is unavailable during the SWITCH phase. You specify a new keyword, FASTSWITCH, which preserves the data set name and updates the catalog to refer to the newly reorganized dataset. The time savings can be quite significant when DB2 is reorganizing hundreds of table spaces and index objects.
- Additional parallel processing improves the elapsed time of the BUILD2 phase of REORG SHRLEVEL(CHANGE) or SHRLEVEL(REFERENCE).

More concurrency with online LOAD RESUME

In prior releases of DB2, access to data during LOAD processing was prohibited. With Version 7 you can choose whether users have read and write access to the data during LOAD processing. You can increase availability by loading data concurrently with user transactions.

Changing system parameters without stopping DB2

Support online transaction processing and e-business without interruption. One of the causes for planned outages is the need to change one or more of the subsystem parameters, known as ZPARMS. With Version 7, you can change many ZPARMS without stopping DB2.

Improved availability of some user objects

Version 7 gives you more flexibility to control the availability of your objects that are associated with failed or canceled transactions. You can cancel a thread and specify that no rollback processing is to be performed. DB2 removes some restrictions imposed by the RESTART function and adds a NOBACKOUT option of the CANCEL THREAD command.

Improved availability with CHECKPAGE for COPY

DB2 for OS/390 and z/OS can perform enhanced error-checking on data pages and ensure that your backup copy is error-free. You can choose to have the COPY utility check both data pages and index pages.

Log suspend and resume for external copy

The LOG SUSPEND command suspends update activity and logging while you make an external copy of your production system. The LOG RESUME command restarts update activity and logging. During the brief suspension, you can use a fast-disk copy facility, such as Enterprise Storage Server™ FlashCopy or RAMAC® Virtual Array SnapShot, to make a copy.

Allow retry of log-read request

Using earlier versions of DB2 for OS/390, the DB2 subsystem terminates whenever a log-read request fails during a “must-complete” operation, such as rollback. Sometimes the condition causing the log-read failure is correctable, such as a temporary HSM or tape subsystem problem. With Version 7, you receive ample warning so that you can take corrective action and avoid a DB2 outage. When a log-read failure occurs during a “must-complete” operation, DB2 issues a write to operator with reply (WTOR) so that you can correct the error. You can retry the log read, and avert a DB2 outage.

DDF suspend and resume

You can use DDF SUSPEND commands and DDF RESUME commands at a server to temporarily suspend activity from requesters without terminating connections. By suspending requester activity, you let data definition operations at the server complete.

Index access for small tables

DB2 adds the ability to select index access for small tables. You provide a value on a subsystem parameter that determines the threshold for a small table.

Data sharing enhancements

DB2 for OS/390 and z/OS improves availability for data sharing environments in Version 7.

Restart light

A new feature of the START DB2 command allows you to choose restart light for a DB2 member. Restart light allows a DB2 data sharing member to restart with a minimal storage footprint, and then to terminate normally after DB2 frees retained locks. The reduced storage requirement can make a restart for recovery possible on a system that might not have enough resources to start and stop DB2 in the normal mode. If you experience a system failure in a Parallel Sysplex, the automated restart in *light mode* removes retained locks with minimum disruption.

Consider using DB2 restart light with restart automation software, such as OS/390 Automatic Restart Manager.

Improvements for IMMEDIATE bind option

A Version 6 enhancement offers you the choice to immediately write updated group-buffer-pool dependent buffers. In Version 7, the option is reflected in the DB2 catalog and externalized on the installation panels.

Persistent structure size changes

In earlier releases of DB2, any changes you made to structure sizes using the SETXCF START,ALTER command might have been lost when you rebuilt a structure and recycled DB2. Now you can allow changes in structure size to persist when you rebuild or reallocate a structure.

V6

Recent enhancements for managing a data sharing environment

Several recent enhancements to Version 6 improve your ability to manage your data sharing environment.

Faster shutdown of DB2 data sharing members: You can more easily apply service or change system parameters. A new CASTOUT(NO) option on the -STOP DB2 command enables a faster shutdown of DB2 data sharing members.

New global scope for IFI calls and for options of several commands: Several recent enhancements to the Instrumentation Facility Interface (IFI) and some commands help you manage your data sharing environment more easily. Now, data from all the data sharing members can be available with a single call from any one of the members.

The following commands have new options with a global scope:

- Start Trace
- Display Trace
- Stop Trace
- Display Thread
- Start Procedure
- Display Procedure
- Stop Procedure
- Start Function Specific
- Display Function Specific
- Stop Function Specific

Outstanding performance when running DB2 on Enterprise Storage Server

You can achieve outstanding performance for your queries, utilities, and transactions by installing and running DB2 on IBM Enterprise Storage Server (ESS). Using large cache, fast disks, and parallel access volumes, ESS and DB2 deliver outstanding results for a wide variety of DB2 for OS/390 and z/OS workloads. To learn more about ESS and DB2 performance results, see IBM's Enterprise Storage Server Web site. See "Helpful Web sites" on page 26 for the site address.

Easier integration of enterprise applications

Recent releases of DB2 for OS/390 and z/OS deliver function that helps you to access large data stores, to develop and to maintain client-server applications, and call stored procedures from popular desktop tools. Version 7 enables you to more easily integrate applications that access data in many environments.

DB2 XML Extender for OS/390 and z/OS for your e-business

DB2 XML Extender for OS/390 and z/OS is a new member of the DB2 Extender family. You can use DB2 XML Extender to store, retrieve, and search XML documents in a DB2 database. DB2 XML Extender for OS/390 and z/OS supplies the data types, functions, and stored procedures to let you manage XML documents.

You can use DB2 XML Extender for OS/390 and z/OS to:

- Store XML documents as whole documents in columns.
- Compose XML documents from a collection of data stored in multiple columns and tables.
- Decompose an XML document back into data stored in columns and tables based on the previously defined mapping.

An application server can send the XML documents over the Internet to other sites.

XML is the standard for data interchange for the next generation of electronic business-to-business and business-integration solutions. Use interchange formats that are based on XML to leverage your critical business information in DB2 databases in business-to-business solutions. When you store, retrieve, and search XML documents in a DB2 database, you benefit from the unmatched performance, reliability and scalability of DB2 for OS/390 and z/OS. With DB2 XML Extender for OS/390 and z/OS you can integrate Internet applications, that are based on XML documents, with your existing DB2 database.

Storing, transmitting, and searching XML documents

DB2 XML Extender for OS/390 and z/OS supports storing XML documents in DB2 databases in two different ways:

Collections: You can store XML documents as untagged and traditional data items in multiple columns and tables in the DB2 database. You can also create XML documents from your existing data, as shown in Figure 5 on page 17.




Figure 5. An XML document stored as a collection of traditional data items in multiple columns and tables in DB2

You define a mapping of the document elements to the data in the database. You can then use DB2 XML Extender's stored procedures to compose XML documents from data that is stored in DB2 tables, based on the previously defined mapping. You can send the resulting XML document over the Internet, by way of an application server, to a corresponding application server at another site. You can also receive XML documents and decompose them into data that is stored in the DB2 tables based on these mappings.

Columns: You can store XML documents in a DB2 database as a whole document in a single column with a new XML data type, as shown in Figure 6.




Figure 6. An XML document stored in a single column of a DB2 table with a new XML data type

You use SQL to do fast and powerful searches on XML distinct types just as you would search other distinct types. You can automatically extract elements or attributes into traditional SQL data types. Then you can use the sophisticated indexing and SQL query capabilities of DB2 to search the data. In addition, DB2 Text Extender for OS/390 and z/OS supports structured documents such as XML. You can apply the powerful text search functions in DB2 to a section or list of sections within a set of XML documents. This can significantly improve the effectiveness of the search.

Graphical support for XML Extender

DB2 XML Extender for OS/390 and z/OS provides a wizard to easily perform administration tasks. The extender helps you enable your database for XML documents, store DTDs, and map XML elements and attributes to DB2 tables and columns.

Unions in views

Version 7 adds support for the UNION and UNION ALL operators in a view definition, an insert, an update, a nested table expression, or a subquery predicate. This enhancement increases programmer productivity by allowing complex unions in views, rather than in each cursor. It also makes porting SQL easier.

DB2 can treat a series of tables as if the series were a single table by allowing you to create a view using UNION and UNION ALL operators. DB2 attempts to optimize a query that references such a view by:

- Distributing the joins, predicates, and aggregation in the query to the subselects of UNION ALL when possible
- Avoiding materialization when possible
- Eliminating subselects if the query does not need them

The INSERT statement now supports fullselect rather than subselect and the UPDATE statement supports row-fullselect rather than row-select.

This function improves DB2 family compatibility and complies with SQL99 standards.

More flexibility with SQL

You have more choices when using SQL and greater compatibility with DB2 on other operating systems.

Scrollable cursors

Support for scrollable cursors lets applications use a powerful new set of SQL statements to fetch data using a cursor at random and in forward and backward directions. The syntax can replace cumbersome logic techniques and coding techniques and improve performance. Scrollable cursors are especially useful for screen-based applications. You can specify that the data in the result table remain static or do the data updates dynamically. You can specify that the data in the result table remain insensitive or sensitive to concurrent changes in the database. You can also update the database if you choose to be sensitive to changes. For example, an accounting application can require that data remain constant, while an airline reservation system application must display the latest flight availability information.

Self-referencing subselect on UPDATE or DELETE

In previous releases of DB2, in a searched UPDATE and DELETE statement, the WHERE clause can not refer to the object being modified by the statement. Version 7 removes the restriction for the searched UPDATE and DELETE statements, but not for the positioned UPDATE and DELETE statements. The search condition in the WHERE clause can include a subquery in which the base object of both the subquery and the searched UPDATE or DELETE statement are the same.

In Figure 7 on page 19, the application gives a 10-percent increase to each employee whose salary is below the average salary for the job code. The base object for both the UPDATE statement and the WHERE clause is the EMP table. DB2 evaluates the complete subquery before performing the update.

```
UPDATE EMP X SET SALARY = SALARY * 1.10
WHERE SALARY < (SELECT AVG(SALARY) FROM EMP Y
WHERE X.JOBCODE = Y.JOBCODE);
```

Figure 7. Updating the EMP table using a self-reference subselect.

V6

UPDATE with subselect

You can use a subselect to determine the value that is to be used in the SET clause of an UPDATE statement. This enhancement improves DB2 family compatibility by delivering DB2 function that is available in other environments.

FETCH FIRST *n* ROWS SQL clause and fast implicit close

A new SQL clause and a fast close improve performance of applications in a distributed environment. You can use the FETCH FIRST *n* ROWS clause to limit the number of rows that are prefetched and returned by the SELECT statement. You can specify the FETCH FIRST ROW ONLY clause on a SELECT INTO statement when the query can return more than one row in the answer set. This tells DB2 that you are only interested in the first row, and you want DB2 to ignore the other rows.

Fast implicit close means that during a distributed query, the DB2 server automatically closes the cursor after it prefetches the *n*th row if you specify FETCH FIRST *n* ROWS ONLY, or when there are no more rows to return. Fast implicit close can improve performance because it saves an additional network transmission between the client and server.

Support for USER and USING keywords on the CONNECT statement

For DB2 family compatibility, IBM adds support for USER and USING keywords to the CONNECT statement. You can easily port your applications developed on the workstation to DB2 for OS/390 and z/OS. The enhancements also allow Websphere to reuse DB2 connections for different users and enable DB2 for OS/390 and z/OS to check passwords.

ORDER BY *expression*

A new option, ORDER BY *expression*, lets you specify an expression with operators as the sort key for the result table of a SELECT statement (that is, not simply a column name or integer). The query to which ORDER BY is applied must be a subselect to use this form of the sort key for the ORDER BY clause.

SQL scalar functions

DB2 introduces SQL scalar functions, which provide a fast and easy way to write simple user-defined functions. You write code for an SQL scalar function when you define it, which eliminates the need to write and prepare a host-language program. In addition, because the source code for an SQL scalar function is stored in the DB2 catalog, an SQL scalar function performs better than an external user-defined function that provides the same function.

Global transaction support for distributed applications

Distributed applications can take advantage of global transaction support. Independent DB2 agents can share locks and access the same data. One or more of those DB2 agents can perform update operations. Your transaction manager must coordinate commit operations among the various DB2 agents using a two-phase commit protocol.

Better monitoring of the distributed environment

You can use the new `-DISPLAY DDF` command to obtain information about the status of DDF, for example, whether it is stopped, started, or suspended. You can also use the `DETAIL` keyword to obtain information such as the maximum number of DBATs allowed, the maximum number of connections allowed, or the number of queued threads.

Support for loading output from a `SELECT` Statement

A new extension to the `LOAD` utility provides synergy between the power of the `LOAD` utility and the connectivity function and reliability of `DRDA`. Now, you can load the output of any `SELECT` statement directly into a table on DB2 for OS/390 and z/OS. Since the `SELECT` statement can access any `DRDA` server, the data source can be any member of the DB2 family, DataJoiner®, or any other vendor that supports `DRDA` server capabilities.

This function is much simpler and easier to use than the process of unloading data, transferring the output file to the target site, and then running the `LOAD` utility. Use of this function also avoids the file size limitations on some operating systems.

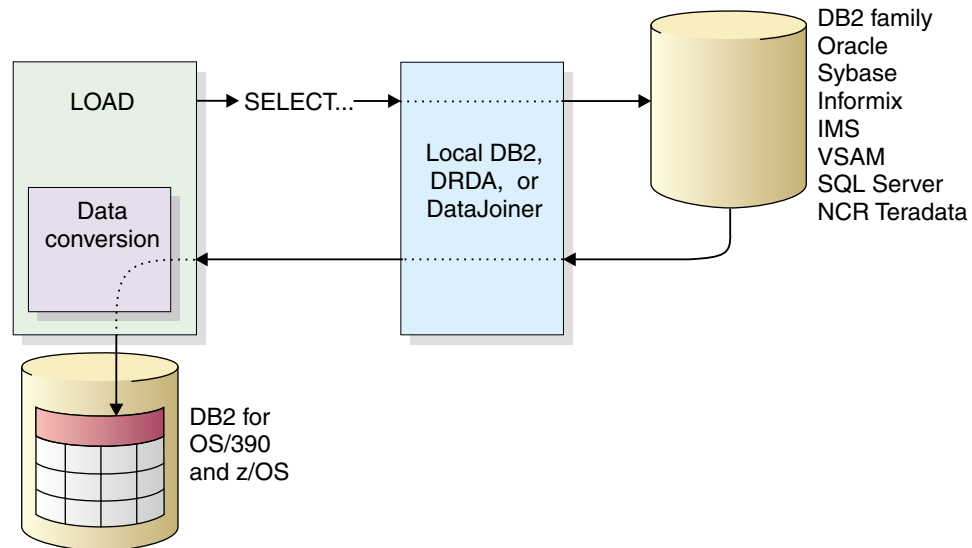


Figure 8. Loading tables with output from a `SELECT` statement

Support for Java stored procedures

DB2 for OS/390 and z/OS now handles stored procedures that are written in the Java programming language. You can write Java stored procedures that contain either static SQL (by using `SQLJ`) or dynamic SQL (by using `JDBC`). Alternatively, you can use the DB2 Stored Procedures Builder (an element of the DB2 Management Clients Package) to build Java stored procedures. Java stored

procedures that run on DB2 for OS/390 and z/OS can also run on other operating systems and platforms, including those of other database vendors, without being recompiled or modified.

Version 6 of DB2 for OS/390 runs compiled Java stored procedures under Visual Age for Java. Version 7 of DB2 for OS/390 and z/OS processes interpreted Java stored procedures that run in a Java virtual machine.

V6

Stored procedures that are written entirely in SQL

With SQL procedures, you can now write stored procedures that consist entirely of SQL statements. An SQL procedure can include declarations of variables, conditions, cursors, and handlers. The SQL procedure can also include flow control, assignment statements, and traditional SQL for defining and manipulating relational data. These extensions provide a procedural language for writing stored procedures, and they are consistent with the Persistent Stored Modules portion of the SQL standard. You can use DB2 Stored Procedures Builder to build SQL stored procedures. For more information about DB2 Stored Procedures Builder, see “Building DB2 stored procedures from your workstation” on page 37.

ODBC 3.0 support

DB2 for OS/390 and z/OS provides ODBC 3.0. support for many ODBC 3.0 APIs. With this added function, DB2’s ODBC support aligns more closely with industry standards, enhances DB2 family compatibility, and increases application portability.

Savepoints

Your application can set a savepoint within a transaction. Application logic can undo the data and schema changes that were made since the application set the savepoint without affecting the overall outcome of the transaction. Using savepoints makes coding applications more efficient. You no longer need to include contingency and what-if logic in your applications.

Support for defining identity columns

An *identity column* in a table has an attribute that enables DB2 to automatically generate a unique numeric value for each row that is inserted into the table. Identity columns are ideal for generating unique primary key values. Applications that use identity columns might be able to avoid concurrency and performance problems that sometimes occur when applications implement their own unique counters.

Ability to define declared temporary tables

Declared temporary tables complement the existing created temporary tables that are available in Version 5 of DB2 for OS/390 and z/OS. Declared temporary tables do not have descriptions in the catalog tables. The tables support indexes, UPDATE statements, and DELETE statements. You can implicitly define the columns and use the result table from a SELECT.

SQL support for REXX

With the REXX-support enhancement, application programmers can issue SQL statements from REXX programs. The SQL statements can be anywhere a REXX command can be. Programmers can also write DB2 stored procedures in REXX. The SQL interface to REXX supports almost all SQL statements that DB2 for

OS/390 and z/OS supports. REXX applications can be executed in all REXX environments, including UNIX System Services.



New Precompiler Services

In DB2 Version 7 compiler writers can modify their compilers to invoke the new component, Precompiler Services, to produce an *SQL statement coprocessor*. An SQL statement coprocessor performs the same functions as the DB2 precompiler, but it performs those functions at compile time. If your compiler has an SQL statement coprocessor, you can eliminate the precompile step in your batch program preparation jobs.

Using an SQL statement coprocessor has several advantages over using the DB2 precompiler. The SQL statement coprocessor:

- Removes restrictions on DB2 host variables
- Makes debugging a program easier
- Removes special restrictions on the format of the source file (other than the restrictions of the compiler)

IBM COBOL for OS/390 & VM Version 2 Release 2 has an SQL statement coprocessor that is called the COBOL DB2 coprocessor. The DB2 precompiler has some restrictions that the COBOL DB2 coprocessor does not have. When you process SQL statements with the COBOL DB2 coprocessor, your program can include:

- Fully-qualified names for structured host variables
- SQL statements at any level of a nested COBOL program, instead of in the top-level source file only
- Nested SQL INCLUDE statements
- COBOL REPLACE statements to replace text strings in SQL statements

Using the host language compiler that supports Precompiler Services enhances DB2 family compatibility. You can more easily import applications from other database management systems and from other operating environments into DB2 for OS/390 and z/OS.

Unicode support for international data stores and e-business

Multinational companies that engage in international trade often store data from more than one country. Some countries use different coded character set identifiers. Previous releases of DB2 have offered support for numerous code sets of data in either ASCII or EBCDIC format. However, there was a limitation of one code set per system.

Version 7 of DB2 for OS/390 and z/OS delivers support for Unicode encoded data. The encoding scheme can represent the code points of many different geographies and languages. You can easily store multilingual data within the same table or on the same DB2 subsystem.

Enhanced DB2 Extender for text objects

DB2 for OS/390 Version 6 introduced DB2 Extenders™. Using DB2 Extenders, you can store and manipulate image, audio, video, and text objects. With DB2 Version

7, you can install and begin using the extender for text objects more quickly. If you install DB2 on Version 2 Release 8 of OS/390, which has an integrated search capability for text objects, you need not download the text extenders package from the Web. DB2 improves documentation to include better guidance through the installation steps for the extender. Performance is improved, and you have more choices to customize the extender.

Enhanced business partnerships with independent software developers

In today's competitive business environment, organizations must build on a solid database foundation to deploy enterprise applications. You need applications that can evolve to meet the changing needs of your business. You can develop these applications yourself, or you can purchase and implement applications from a variety of software companies. If your enterprise data is in DB2, however, you want these applications to take advantage of the scalability, reliability, and industrial-strength performance of DB2.

Working closely with IBM, the leading software developers create enterprise solutions that are designed to leverage the reliability and industrial-strength performance of DB2. Several software developers (such as Baan, SAP, Siebel Systems, PeopleSoft, Tivoli®, Vantive, MicroStrategy, and Forte) work closely with the IBM DB2 development team to identify new functions and capabilities for DB2 users. IBM works with these independent software development companies to understand the technical requirements of their applications. Many of the functions described in this book are the result of these partnerships, including Unicode support, utility lists with pattern matching and dynamic allocation, and UNION IN VIEW.

Improvements in connectivity

IBM improves e-business connectivity with new levels of support for both ODBC and JDBC, the ability to perform commit and rollback function in a stored procedure, Kerberos security for your clients, and the server elapsed time reporting.

Support for COMMIT and ROLLBACK in stored procedures

Earlier releases of DB2 for OS/390 and z/OS do not allow a COMMIT or a ROLLBACK from within a stored procedure. Version 7 removes that restriction in many cases. DB2 commits or rolls back the whole unit of work, including uncommitted changes that are made from the calling application, before the stored procedure call.

Support for Windows Kerberos security

You can more easily manage your workstation clients who seek access to data and services from heterogeneous environments. DB2 support for Kerberos authentication:

- Eliminates the flow of unencrypted user IDs and passwords across the network.
- Enables single-logon capability for DRDA® clients by using the Kerberos principle name as the global identity for the end user.
- Simplifies security administration by using the Kerberos principle name for connection processing and by automatically mapping the name to the local user ID.

- Uses the Resource Access Control Facility (RACF®) product to perform much of the Kerberos configuration. RACF is a familiar environment to administrators of OS/390.
- Eliminates the need to manage authentication in two places, the RACF database, and a separate Kerberos registry.

Reporting server elapsed time at the workstation

Workstations accessing DB2 data can now request that DB2 return the elapsed time of the server, which is used to process a request in reply from the DB2 subsystem. The server elapsed time allows remote clients to quickly determine the amount of time it takes for DB2 to process a request. The server elapsed time does not include network delay time. As a result, workstation clients can determine, in real time, the performance bottlenecks among the client, the network, and DB2.

New packaging for DB2 utilities

DB2 for OS/390 and z/OS Version 7 delivers utilities in various packages. Core utilities are included with DB2. Other utilities are available in three separate products. The core utilities include:

- CATMAINT
- DIAGNOSE
- LISTDEF
- OPTIONS
- QUIESCE
- REPAIR
- REPORT
- TEMPLATE
- All standalone utilities

The DB2 utility products are:

- DB2 Diagnostic and Recovery Utilities (5655-E62)
 - CHECK DATA
 - CHECK INDEX
 - CHECK LOB
 - COPY
 - COPYTOCOPY
 - MERGECOPY
 - MODIFY RECOVERY
 - MODIFY STATISTICS
 - REBUILD INDEX
 - RECOVER
- DB2 Operational Utilities (5655-E63)
 - COPY
 - EXEC SQL
 - LOAD
 - REBUILD INDEX
 - RECOVER
 - REORG INDEX
 - REORG TABLESPACE
 - RUNSTATS
 - STOSPACE
 - UNLOAD
- DB2 Utilities Suite (5697-E98)
 - CHECK DATA
 - CHECK INDEX

- CHECK LOB
- COPY
- COPYTOCOPY
- EXEC SQL
- LOAD
- MERGECOPY
- MODIFY RECOVERY
- MODIFY STATISTICS
- REBUILD INDEX
- RECOVER
- REORG INDEX
- REORG TABLESPACE
- RUNSTATS
- STOSPACE
- UNLOAD

All DB2 utilities operate on catalog, directory and sample objects and require no additional products.

Support for Migration to Version 7 of DB2 for OS/390 and z/OS from Version 5 or Version 6

Version 7 of DB2 for OS/390 and z/OS supports migration from DB2 for OS/390 Version 5 or DB2 for OS/390 Version 6. To plan for your migration, consult the following publications:

- DB2 Release Planning Guide for prerequisites of the product
- DB2 Installation Guide for migration considerations

Both publications are available in BookManager® and PDF formats on the Web at the following location:

- ibm.com/software/db2os390/library.html

Information to help you use DB2

IBM delivers some new information for DB2 for OS/390 and z/OS users and frequently updates the DB2 Web sites for easy reference.

New books about DB2 for OS/390 and z/OS

The DB2 for OS/390 and z/OS Version 7 library is designed to help your perform the various tasks associated with a database management system. The library includes two new books:

An Introduction to DB2 for OS/390

This book helps new users learn about DB2 very quickly. This book:

- Provides a comprehensive overview of DB2 for OS/390 and z/OS.
- Explains basic concepts associated with relational database management systems.
- Gives new users what they need to know before they begin using the current version of the product.

DB2 XML Extender for OS/390 and z/OS Administration and Programming

Database administrators, system programmers, XML developers, and document application developers are the audience for this book. They can obtain guidance for developing, deploying, and managing databases and database applications that store XML documents from traditional data.

This book describes how to use and manage the DB2 XML Extender for OS/390 and z/OS and includes the following topics:

- How to choose a methodology for storing and accessing XML documents
- How to prepare a database server, tables, and columns for DB2 XML Extender for OS/390 and z/OS for either column storage or collection storage
- How to create and use text indexes
- How to use DB2 XML Extender for OS/390 and z/OS user-defined types, user-defined functions, and stored procedures

You can order publications for DB2 for OS/390 and z/OS Version 7 when the product becomes generally available. IBM provides books on CD-ROM in BookManager format for online viewing and the Portable Document Format (PDF) for online viewing and printing. IBM also installs both formats of the books on the DB2 for OS/390 and z/OS Web site.

IBM optimizes the PDF files so you can easily retrieve information online. Use the handy table of contents, which is always visible in a frame, to navigate to information that you require. Print just a few pages of the book when you need the information, or print the entire book for desk reference. IBM periodically refreshes information on CD-ROM and the Web site.

DB2 for OS/390 and z/OS on the Web

The DB2 for OS/390 and z/OS Web site delivers the latest news about DB2 for OS/390 and z/OS , hints and tips for using the product, some software that you can download, and links to many related sites. The site has many DB2 for OS/390 and z/OS books in PDF and BookManager formats that you can download without paying a fee to IBM.

Helpful Web sites

For more information about some of the topics discussed in this book, go to the following Web sites.

DB2 family

DB2 Universal Database for OS/390 and z/OS

ibm.com/software/db2os390

DB2 Universal Database for UNIX, Windows, OS/2

ibm.com/software/data/db2/udb

DB2 Universal Database for AS/400

ibm.com/as400/db2

DB2 Server for VSE and VM

ibm.com/software/data/db2/vse-vm

Data Management Tools

ibm.com/software/data/db2imstools/

DB2 for OS/390 and z/OS performance on Enterprise Storage Servers

ibm.com/software/db2os390/pdf/db2ess.pdf

Case studies of successful e-business implementations using DB2

ibm.com/software/casestudies

ibm.com/software/data/db2/os390/casestudies

DB2 family software partners

ibm.com/software/data/partners

ibm.com/solutions/businessintelligence/partners/index.htm

DB2 Magazine

db2mag.com

Other references**Application development**

ibm.com/software/ad

Business intelligence and data warehouse

ibm.com/solutions/bi

Database and data management

ibm.com/software/data

DeveloperWorks

ibm.com/developer

e-business

ibm.com/ebusiness

Enterprise Storage Servers

ibm.com/storage

IBM Redbooks

ibm.com/redbooks

Java technology

ibm.com/java

Subscription service for announcements about IBM products

ibm.com/isource

XML for developers

ibm.com/XML

Chapter 3. Managing your Enterprise

Version 7 of DB2 for OS/390 and z/OS offers several features that help you integrate, analyze, summarize, and share data across your enterprise. IBM adds a new extender for searching your Web applications and a new feature for building a data warehouse. IBM has strengthened its focus on data management tools. The following sections describe the features and tools:

“Integrating and analyzing your business information using the DB2 Warehouse Manager feature”

“Managing your enterprise with the DB2 Management Clients Package” on page 34

“Web-ready applications” on page 42

“Data Management Tools for your database” on page 44

Integrating and analyzing your business information using the DB2 Warehouse Manager feature

The DB2 Warehouse Manager feature brings together the tools to build, manage, govern, and access DB2 for OS/390 and z/OS-based data warehouses. The DB2 Warehouse Manager feature uses proven technologies with new enhancements that are not available in previous releases. The DB2 Warehouse Manager feature delivers tightly integrated components that enable you to do the following tasks:

- Simplify prototyping, development, and deployment of your warehouse
- Give control to your data center to govern queries, analyze costs, manage resources, and track usage
- Help your users find, understand, and access information
- Give you more flexibility in the tools and techniques you use to build, manage, and access the warehouse
- Meet the most common reporting needs for enterprises of any size

DB2 Warehouse Manager

The DB2 Warehouse Manager feature for OS/390 provides the following components:

- A restricted license for Version 7 of DB2 Universal Database Enterprise Edition, which delivers the database warehouse management infrastructure and OLAP Starter Kit. This infrastructure replaces Visual Warehouse with enhancements that include:
 - The Data Warehouse Center, as shown in Figure 9 on page 30, a new graphical interface that is integrated with the DB2 Control Center. The Data Warehouse Center which provides the administration for building and managing data warehouses. New features include:
 - A navigator window to display and browse warehouse objects
 - A schema modeler that visualizes table relationships
 - The ability to extend the environment with user-written stored procedures or user-defined functions
 - A process modeler that provides a canvas and palette for defining warehouse steps and control flow between them
 - Integration with the DB2 OLAP Integration Server
 - Integration with DB2 DataPropagator to define subscriptions for warehouse population

- An integrated tutorial for fast learning
- A warehouse manager, which runs on Microsoft Windows NT and controls all the operations that are defined through the Data Warehouse Center.
- A warehouse agent for Windows NT, which executes local operations on behalf of the components of the DB2 Warehouse Manager feature.
- The OLAP Starter Kit, which is a limited user license for DB2 OLAP Server® Integration Server for building and deploying OLAP applications.

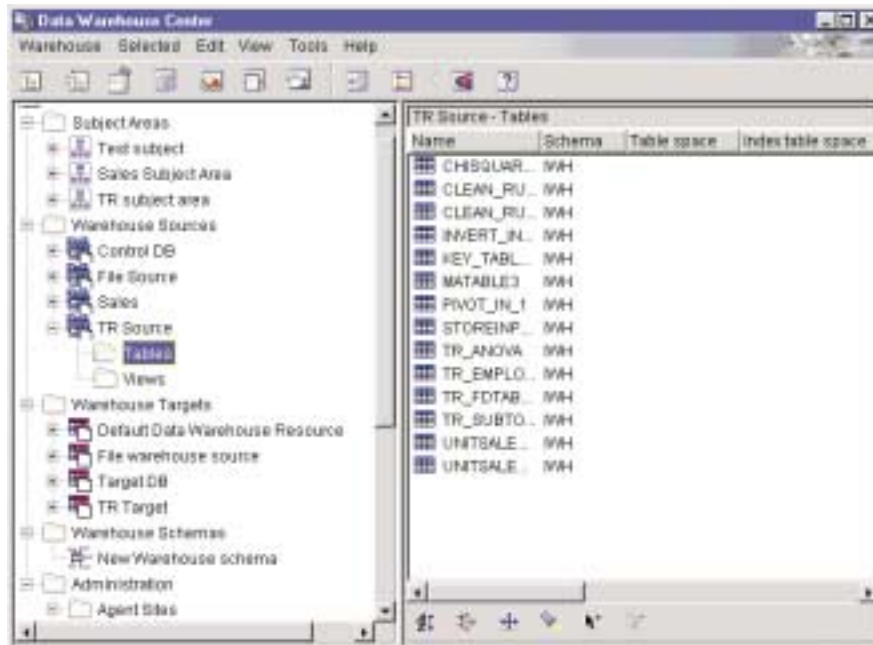


Figure 9. DB2 Data Warehouse Center

- Warehouse agents for OS/390, UNIX, AS/400, and Windows NT which execute processes on behalf of the components of the DB2 Warehouse Manager feature.
- Prebuilt programs for performing a variety of OS/390 specific tasks, such as executing file transfer protocol (FTP), submitting job control language (JCL), and triggering server or client programs.
- Warehouse transformers for OS/390, which are stored procedures or user-defined functions. These transformers perform complex transformations that are commonly used in warehouse development including data manipulation, data cleansing, key generation, and statistical calculations. These transformers augment the transformations that are available with built-in SQL, including SQL string manipulation, Boolean operations, mathematical calculations, data type conversions, and utility invocations.
- The enhanced Information Catalog Manager, previously available with Visual Warehouse, improves the Web interface. The Information Catalog Manager provides tighter integration with DB2 OLAP Server, QMF™ from any DB2 database, Hyperion's Essbase, other front-end tools, and IBM Enterprise Information Portal (EIP).
- The QMF family of products:
 - QMF for OS/390, which enables your users to issue queries on OS/390
 - QMF High Performance Option (HPO)

- QMF for Windows, which helps you leverage your QMF for OS/390 investment by reusing query and report definitions and extending QMF to the Windows environment for integration with desktop tools

Reporting and governing your enterprise using QMF

Query Management Facility (QMF) is the tightly integrated, powerful, and reliable tool for query and reporting within IBM's DB2 family. QMF provides new capabilities that improve user productivity and enhance functionality. With QMF, you can work with any DB2 data from DB2 for OS/390 and z/OS to DB2 for VSE and VM, to DB2 for AS/400 to DB2 on workstation servers that run OS/2, Windows NT, AIX and other UNIX operating systems. QMF supports DB2 on large parallel processors. When coupled with DB2 DataJoiner, QMF allows access to nonrelational and other vendor data sources.

QMF is a family of integrated tools that offers a total solution for your enterprise query and reporting needs. With these tools, you can access large amounts of data, share central repositories of queries and reports, and publish reports to the Web for viewing with a browser. You can also allow a broader set of users to access QMF, knowing that your applications are secure.

QMF Version 7

QMF Version 7 enhancements include support for the following functions:

- New default edit codes for current DB2 date (DD) and time (TT) formats in date columns and time columns
- Direct navigation to the QMF HOME panel with the SHOW command
- Comprehensive defaults for object types in many commands, including:
 - CONVERT
 - EDIT
 - EXPORT
 - PRINT
 - RESET
 - RUN
 - SAVE
- Functions that are easier to use with prompting for commands and defaults that are preloaded in prompts so users do less typing
- Extended text option in the MSG command, eliminating the need to add double quotation marks when the text of a message has single quotes
- Two DB2 data types: ROWID and limited LOB support
- Remote-unit-of-work application request in DB2 for VSE using Customer Information Control System (CICS)
- Server support for DB2 for AS/400 Version 4 Release 4 or later releases
- Installation improvements so that you can use one QMF install base that you bind on other platforms

QMF High Performance Option

QMF HPO provides a set of integrated tools for performance management within the QMF environment and facilitates the administration of QMF. QMF HPO includes the QMF HPO/Manager and the QMF HPO/Compiler. QMF HPO/Manager consists of a group of facilities that improve governing and object management capabilities. The program includes a preemptive governor to analyze QMF queries. The governing capabilities enable you to establish controls that protect production applications while delivering on-demand information. QMF HPO/Manager provides:

- Scheduling for different user groups by time of day and day of week

- Control over resource consumption, including by numbers of rows and bytes fetched
- The ability to allow and disallow SQL verbs and QMF commands
- Extensive object tracking to gain knowledge about your QMF environment so you can train users and improve performance

QMF HPO/Compiler automatically converts queries and reports into efficient programs in COBOL. The converted programs use static SQL in place of dynamic SQL, which reduces processor usage, DB2 catalog contention, DB2 optimizer overhead, and security concerns.

Version 7 of QMF HPO includes the following enhancements:

- Improved tracking of QMF objects to support very large numbers of QMF users simultaneously
- Preemptive governing to guard QMF workloads from runaway, adhoc queries
- Preemptive governing of QMF batch processes using QMF HPO/Manager
- Ability to bind programs that are generated by QMF HPO Compiler into packages, in addition to or instead of, plans

QMF for Windows

For customers with DB2 databases of many sizes, QMF for Windows provides a query tool with standard Windows capabilities. QMF for Windows delivers several benefits in one package, which includes an intuitive quick-start user interface. You can automate tasks and develop powerful native Windows applications. QMF for Windows provides the following features:

- A full-window table editor that you can use to update DB2 data
- Centralized control over resource use
- Web publishing capabilities and Java-based query and Web publishing capabilities
- A powerful WinSock application programming interface (API) to automate database query, update, and report distribution tasks

QMF for Windows is based on Distributed Relational Database Architecture (DRDA). QMF for Windows provides access to DB2 database servers on all platforms through TCP/IP, Systems Network Architecture (SNA), or Call Level Interface (CLI), as Figure 10 on page 33 shows. To provide this access, QMF uses direct DRDA connectivity without database gateways, middleware, or Open Database Connectivity (ODBC) drivers.

Enhancements to QMF for Windows Version 7 include:

- An improved user interface:
 - QMF for Windows provides a point-and-click and a drag-and-drop interface to QMF Form creation.
You can automatically populate a QMF Form definition by using aggregation, grouping, and formatting.
 - You can convert the resulting display options to QMF form definitions.
 - You can more intuitively navigate and locate QMF objects with a Windows-style tree control.
 - You can create job schedules by calendar, time of day, and more from within QMF for Windows.
 - Java servlet support enables you to perform a Java-based query from a browser.

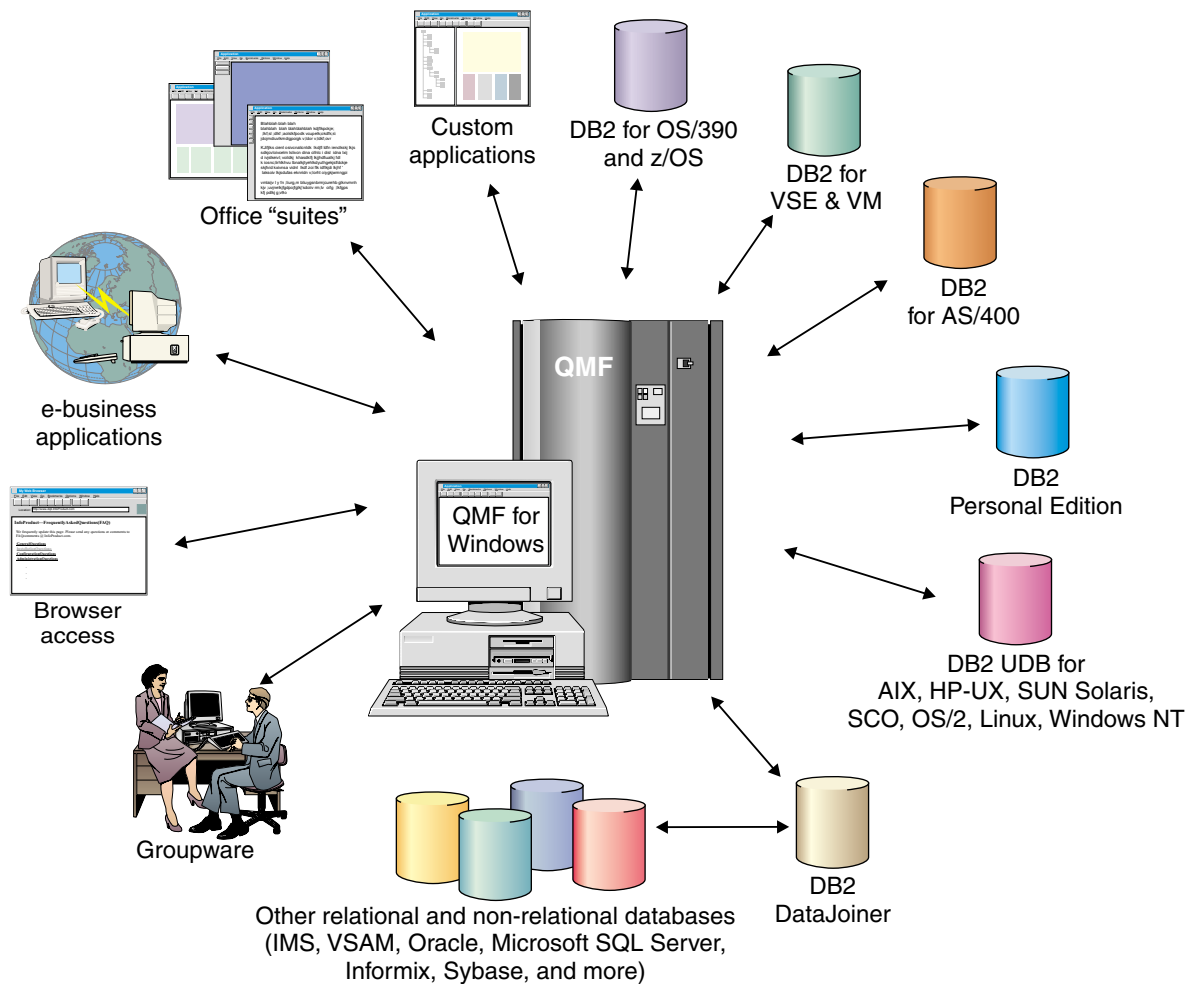


Figure 10. Accessing data using QMF for Windows

- QMF for Windows includes QMF Report Center as an optional install. QMF Report Center was previously a separate download named "Personal Portal,".
- When you invoke the QMF for Windows Table Editor, you can use the DB2 Forms option if you have it installed.
- More command icons are available so that you can customize the toolbar.
- QMF for Windows can remember DB2 passwords when that are used in the Windows NT environment.
- You can sort query results on more than one noncontiguous column and in day-of-week or month-of-year order.
- The Visual Warehouse object library now integrates QMF for Windows objects
- Enhanced database connectivity for distributed applications:
 - Support for stored procedures that return multiple result sets by using the CLI interface
 - Large object support
 - Enhanced support for DB2 for OS/390 and z/OS data sharing
 - Display of CLI-specific information when a CLI connection to the server exists

- New trace options include CLI, TCP/IP, embedded SQL, SQLAM, REXX, CPI-C, and DDM
- More support for QMF for OS/390:
 - Support for all types of QMF for OS/390 procedures, including procedures with REXX logic
 - A new docking toolbar for entering and issuing single QMF commands
- Several new QMF commands:
 - Form conditions and form column definitions
 - A RESET GLOBAL command
 - PF keys that are defined to match QMF for OS/390
- A user interface with national language support for:
 - 19 language translations, which in Version 7 include Belgian French, Canadian French, Swiss French, Swiss Italian, Arabic, and Traditional Chinese
 - Unicode

Managing your enterprise with the DB2 Management Clients Package

The DB2 Management Clients Package is a no-charge feature of DB2 for OS/390 and z/OS Version 7. The Management Clients Package is a collection of workstation-based tools that you can use to work with and manage your DB2 for OS/390 and z/OS environment. The elements of the DB2 Management Clients Package are:

- DB2 Control Center
- DB2 Stored Procedures Builder
- DB2 Installer
- DB2 Visual Explain
- DB2 Estimator

Managing DB2 for OS/390 and z/OS from your workstation with Control Center

The IBM DB2 Universal Database Control Center is a graphical interface that is designed to help you manage DB2 databases on different operating systems. The Control Center can run either as a Java application or as an application on your Web server that your Web browser can access.

With the Control Center for DB2 for OS/390 and z/OS, you can manage DB2 for OS/390 Version 5, Version 6, and Version 7 objects. These objects display on the Control Center main window, along with DB2 Universal Database objects. See Figure 11 on page 35. Using the Control Center, you can create, alter, and drop objects. You can also run utilities that reorganize or load your data in your existing DB2 for OS/390 and z/OS databases.

The Control Center for DB2 for OS/390 and z/OS enhances the usability and administration of DB2 for OS/390 and z/OS objects from the workstation. In Version 7, IBM significantly improves the usability of the DB2 for OS/390 and z/OS utilities.




Figure 11. A DB2 for OS/390 and z/OS subsystem and its objects as displayed in the main navigation window of Control Center

Managing DB2 for OS/390 and z/OS data sets from the Control Center

You can now retrieve a list of data sets that reside on an OS/390 system from the Control Center. The interface to the Multiple Virtual Storage (MVS) file system allows you to work with four types of data sets: A physical sequential data set, a partitioned data set, a partitioned data set extended, or a generation data group. Depending on the data set type, you can rename, delete, or display members of the data set through the Control Center.

Restarting DB2 for OS/390 and z/OS utilities from the Control Center

You can now restart utilities that were originally started from the Control Center. You can execute restart from the last committed point or from the last committed phase. You can define a custom utility identifier from the Control Center, and associate the identifier with any utilities as they run. Customized utility identifiers help you find the utilities that you started among others that are running on the system.

Generating data definitions

You can use the DB2 Administration Tool with Control Center to extract object definitions from the DB2 for OS/390 and z/OS catalog tables. You use the definitions to recreate databases, table spaces, tables, procedures, schemas (for granting privileges), user-defined functions, and distinct types. This feature is useful

for creating a copy of a particular object before making changes to it when fallback might be necessary. You can also use this feature to move or copy an object and its environment to another subsystem. For example, you can copy the table space object, including all its objects, such as tables, indexes, and aliases that are associated with that object, to another subsystem.

Object lists and utility procedures

In the Version 7 Control Center, you can define collections of objects, called object lists. These object lists are combinations of one or more table spaces or index spaces that you can easily generate using pattern-matching characters. With a single command, you can run utilities (like COPY, CONCURRENT COPY, QUIESCE, and REORGANIZE) on every object in the list. By creating a utility procedure, you can define a sequence of object lists or utility combinations that you can run with a single command.

Data set templates

You no longer need to specify a data set name, a data set device, and data set space allocation for each utility. Now you can use data set templates to dynamically allocate data sets that DB2 for OS/390 and z/OS utilities use. Use data set templates to automatically generate the data set naming conventions, the data set parameters, and the media allocation parameters. The default data set templates available with DB2 for OS/390 and z/OS are shown in Figure 12

These are the utilities currently supported by Data set templates.

Users can select from sample templates, or build their own to set as defaults for Data sets when running utilities.

This field displays the Data set naming convention that the template defines.

Utility	Data set description	Template name	Data set name
Concurrent copy	Primary copy data set at local site	CCCOPY	SUSERID.DB2.6SSD.6CB.6SH.D6JDAY.SHOURL6MINUTE
Concurrent copy	Backup copy data set at local site		
Concurrent copy	Filter data set	CCFILTER	SUSERID.DB2.6SSD.6CB.6SH.D6JDAY.SHOURL6MINUTE
Concurrent copy	Primary copy data set at recovery site		
Concurrent copy	Backup copy data set at recovery site		
Copy	Primary copy data set at local site	CCCOPY	SUSERID.DB2.6SSD.6CB.6SH.D6JDAY.SHOURL6MINUTE
Copy	Backup copy data set at local site		
Copy	Primary copy data set at recovery site		
Copy	Backup copy data set at recovery site		
Reorganize index	Unread data set	CCUNLOAD	SUSERID.DB2.6SSD.6CB.6SH.D6JDAY.SHOURL6MINUTE
Reorganize table space	Primary copy data set at local site	CCCOPY	SUSERID.DB2.6SSD.6CB.6SH.D6JDAY.SHOURL6MINUTE
Reorganize table space	Backup copy data set at local site		
Reorganize table space	Discard data set	CCDISCRD	SUSERID.DB2.6SSD.6CB.6SH.DISCARD(+1)
Reorganize table space	Punch data set	CCPUNCH	SUSERID.DB2.6SSD.6CB.6SH.D6JDAY.SHOURL6MINUTE
Reorganize table space	Primary copy data set at recovery site		
Reorganize table space	Backup copy data set at recovery site		
Reorganize table space	Unread data set		
Reorganize table space	Input sort data set	CCSORTIN	SUSERID.DB2.6SSD.6CB.6SH.D6JDAY.SHOURL6MINUTE
Reorganize table space	Output sort data set	CCSORTOUT	SUSERID.DB2.6SSD.6CB.6SH.D6JDAY.SHOURL6MINUTE

Figure 12. The default data set templates available with DB2 for OS/390 and z/OS

In addition to providing standardization, data set templates allow DB2 to calculate the size of output data sets and work data sets. The use of data set templates can reduce utility failures that occur because of miscalculations. In the Control Center, you can associate each utility with default data set templates. You do not need a data set entry when you create utility commands. The initial installation of the DB2 for OS/390 and z/OS enablement for the Control Center includes several templates.

With the default table for the data set templates, you can map the data sets that each utility needs to the default data set templates. The Control Center uses the default table each time you display the utility window. See Figure 12 on page 36 which depicts, the window in the Control Center listing the sample default templates that are installed with DB2 for OS/390 and z/OS. Edit these templates to reflect the conventions that your production subsystem uses.

Information Center

The Control Center includes an Information Center for quick access to DB2 family product information, including:

- Database tasks
- Reference material
- DB2 documentation
- Warehouse administration information
- Troubleshooting aids
- Sample programs for application development on workstations
- DB2 Web-related links

You can use the Information Center to link to current DB2 for OS/390 and z/OS publications in Portable Document Format (PDF) on the Web. You can download the publications to your local server and update the Information Center links to access DB2 for OS/390 and z/OS books directly from your intranet.

Building DB2 stored procedures from your workstation

The IBM DB2 Stored Procedures Builder, an element of the DB2 Management Clients Package, provides an easy-to-use development environment for creating, installing, and testing stored procedures. With the DB2 Stored Procedures Builder, you can focus on creating your stored procedure logic rather than on the details of registering, building, and installing stored procedures on a DB2 server. The DB2 Stored Procedures Builder provides a single development environment that supports the entire DB2 family ranging from the workstation to System/390. Figure 13 depicts the flexibility that is provided by the tool for building and using stored procedures. For example, you can develop stored procedures using IBM Visual Basic™, Microsoft Visual Basic, or Microsoft Visual Studio.

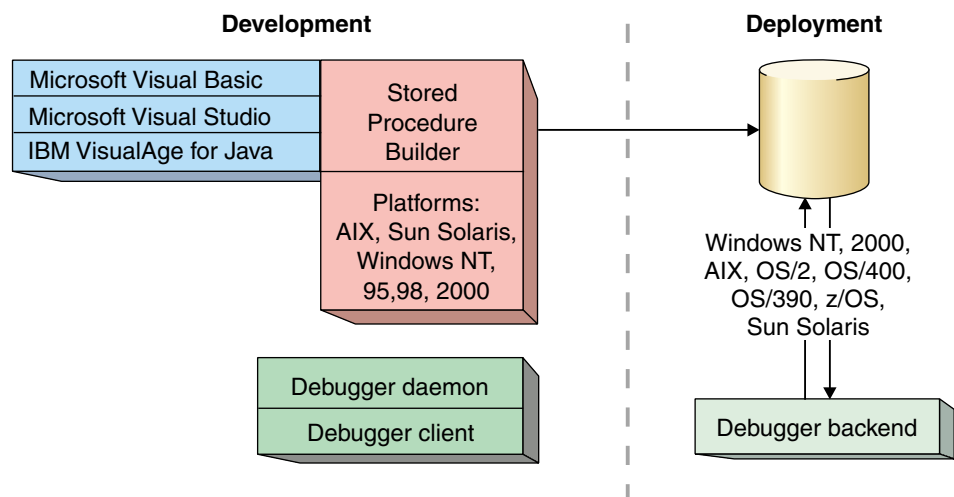


Figure 13. Using DB2 Stored Procedures Builder

With the DB2 Stored Procedures Builder, you can develop stored procedures on one operating system and deploy them on other server operating systems. The DB2

Stored Procedures Builder supports two commonly used languages for stored procedures: SQL procedures language and Java.

The DB2 Stored Procedures Builder has a graphical user interface that guides you through tasks with the help of basic design patterns, SQL assistants, and costing information. Use the DB2 Stored Procedures Builder to perform a variety of tasks that are associated with stored procedures, such as:

- Viewing existing stored procedures
- Modifying existing stored procedures
- Creating new stored procedures
- Running existing stored procedures
- The copying and pasting of stored procedures across connections
- Building stored procedures on target databases in a single step
- Customizing the settings to enable remote debugging of installed stored procedures

The DB2 Stored Procedures Builder requires DB2 Connect. The DB2 Management Clients Package provides a restricted-use copy of DB2 Connect Version 7 to satisfy this functional dependency.

Installing DB2 from a workstation

DB2 Installer is a workstation client that is delivered as an element of the DB2 Management Clients Package. DB2 Installer enhances your productivity whether you are installing DB2 for OS/390 and z/OS for the first time or you are an experienced installer. From your workstation, using a graphical user interface shown in Figure 14 on page 39, you can perform the following tasks:

- Install, migrate, or update DB2 for OS/390 and z/OS from a graphical interface. The graphical interface illustrates the overall installation process and keeps a graphical record of how you defined each subsystem, as shown; see Figure 14 on page 39.
- Customize your DB2 subsystem as much as you need to. You can define a basic subsystem quickly, or you can customize every installation option. The main windows display the parameters that you must specify, and secondary windows display the advanced options.
- Easily control DB2 parameters and run SMP/E, installation, migration, update, fallback, and sample jobs if you have a TCP/IP connection to the DB2 Universal Database Server for OS/390 and z/OS. You receive job status dynamically, and you can edit JCL and examine job output from the workstation.
- Install optional features of DB2 for OS/390 and z/OS, including DB2 Performance Monitor and DB2 DataPropagator.

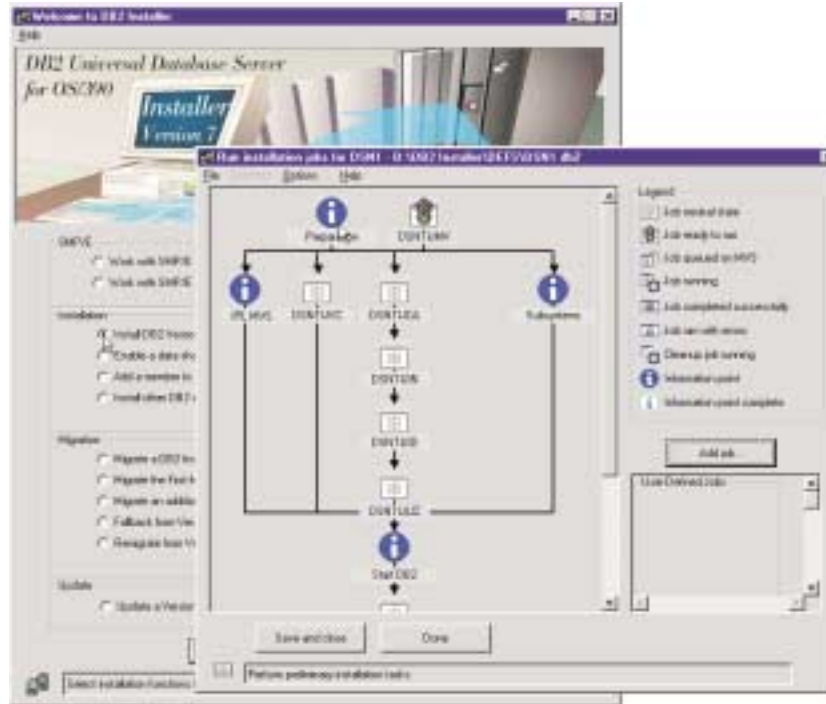


Figure 14. DB2 Installer

Enhancements for Version 7 of DB2 Installer include:

- The ability to use an IBM-supplied stored procedure to download subsystem parameters from your operational DB2 subsystem. You can use the subsystem parameters to migrate or update your subsystem.
- Improved usability of several DB2 Installer windows, including the **File Dialog** windows, the **Before You Begin** window, and the **Welcome** window.
- Support for Windows 95, Windows NT, and Windows 2000.

Using workstation views of DB2 Explain output

DB2 Visual Explain is a workstation client that is an element of the DB2 Management Clients Package. DB2 Visual Explain is an easy-to-use workstation tool that presents the output from DB2 EXPLAIN and dynamic EXPLAIN in a graphical format. Relationships between database objects, such as tables and indexes, are instantly clear as are various operations, such as table space scans and sorts. DB2 Visual Explain also includes a browser for viewing DB2 subsystem parameters. Figure 15 on page 40 displays a graph of an access path.

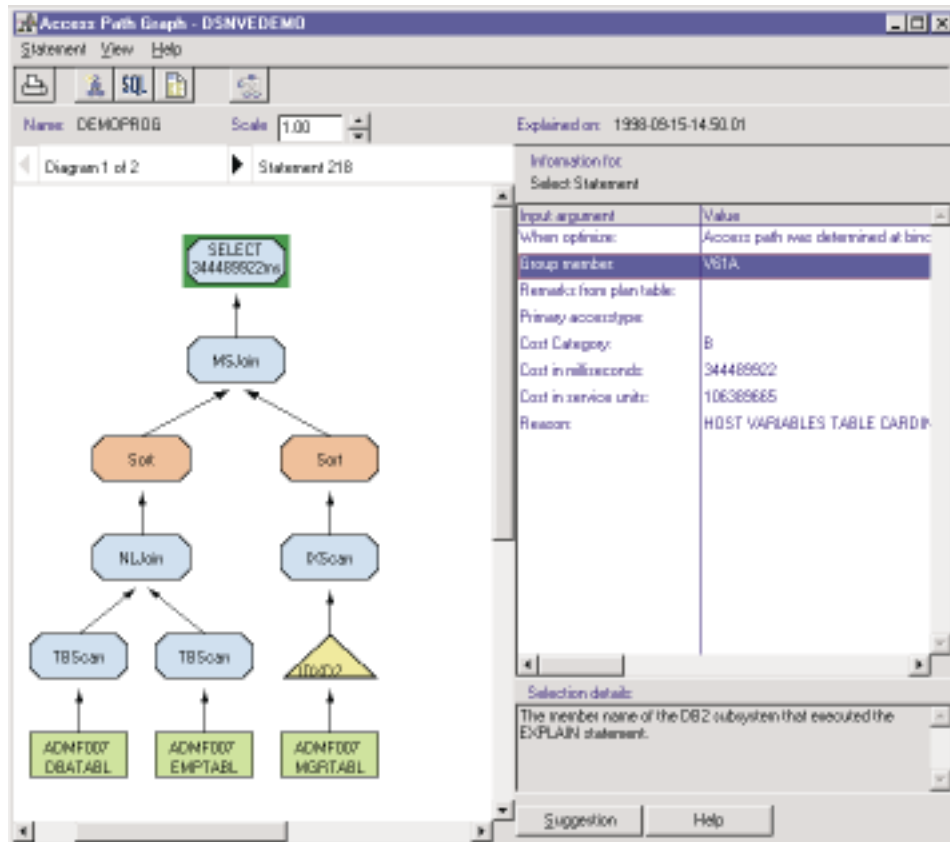


Figure 15. A graph of an access path, displayed by DB2 Visual Explain

If you are using the Control Center, you can start Visual Explain directly from the Control Center.

Use Visual Explain to perform management tasks for DB2 for OS/390 and z/OS:

- View statement costs in milliseconds and service units. The graph of the access path includes the DB2 estimated cost. You can include the cost when viewing explainable statements. Then you can either sort your statements by cost or filter out statements that are based on their costs.
- Filter explainable SQL statements from multiple plans and packages. You can list explainable statements from many plans and packages and then filter them by criteria that you specify. The criteria include statement cost, access path steps (such as table space scans and sorts and either tables or indexes that your SQL statements reference). You can save filter criteria specifications for later use or modification.
- Quickly generate customized reports. The Visual Explain report wizard helps you create customized reports on one or more explainable statements. You choose how detailed a report you want. The report can include the statement cost, a description of the access path, and catalog statistics on the indexes and tables. You can print the report or save it as a text file.
- Specify your own qualifier for catalog tables. You can avoid contention against your DB2 catalog for Visual Explain queries by first copying the catalog tables. Then you can specify your own qualifier to access the copied catalog tables when Visual Explain retrieves DB2 catalog information.

DB2 Visual Explain saves statistics history for comparison with new variations that you enter so that you can improve your access paths.

Estimating DB2 performance with DB2 Estimator

DB2 Estimator is an easy-to-use, standalone tool for estimating the performance of applications for DB2 for OS/390 and z/OS. DB2 Estimator is an element of the DB2 Management Clients Package. Run it on your desktop personal computer, or take it with you on your laptop computer.

With DB2 Estimator, you can model a partial DB2 application or a complete real or planned DB2 application without requiring an actual DB2 system. By providing simple table sizings to a detailed performance analysis of an entire DB2 application, DB2 Estimator saves time, lowers costs, and reduces risk. You can use DB2 Estimator to investigate the impact to your production system of new or modified applications before you implement them. You can do what-if analysis to assess the impact of changes you are considering. Answer many questions, such as:

- What is the impact on your system if the transaction volume doubles?
- What is the impact if your databases increase in size?
- What is the effect on response time if you use a faster processor?
- Is your batch window large enough for your utilities?
- How much storage do I need for the new table and its indexes?

Use DB2 Estimator during all life-cycle phases of a DB2 application. During the initial design phase, you can easily:

- Determine whether your design is optimal and feasible
- Investigate alternative database designs
- Assess the impact of using triggers and different ways of structuring queries and transactions
- Eliminate problems early in the design phase

When creating database objects, use the models that you specified in DB2 Estimator as a guide for naming columns and for specifying attribute values. You can model your system using actual DB2 information by importing information from the DB2 catalog and, if available, from DB2 EXPLAIN or the DB2 Performance Monitor (DB2 PM).

When your application is in production, use DB2 Estimator with tools such as DB2 PM to solve application performance problems. You can evaluate alternative SQL designs, without any risk to your production environment, before you change any production database objects. DB2 Estimator also helps you determine the impact of changes to hardware or workload.

You can use DB2 Estimator for Versions 5, 6, and 7 of DB2 for OS/390 and z/OS. It runs in any environment that supports Windows 95, Windows 98, Windows 2000, or Windows NT. Use DB2 Estimator on any data that you imported from your DB2 for OS/390 and z/OS subsystems. Model an application even when none of the tables, SQL statements, transactions, or configurations exist.

Version 7 of DB2 Estimator adds support for scrollable cursors, `FETCH FIRST n ROWS`, and statistics history.

Web-ready applications

Designed for use with the Web, DB2 supports all the key Internet standards. Built-in Java support, an Extensible Markup Language (XML) Extender, search capability, and a Web server make it easy to deploy e-business applications. Net Search Extender is a new feature and Net.Data for OS/390 is enhanced for Version 7 of DB2 for OS/390 and z/OS.

Text search capability for the Internet with Net Search Extender

Net Search Extender delivers the latest in high-speed search technology. Use its powerful search to provide rapid query responses for your e-business applications. Net Search Extender combines in-memory database technology with text search technology. This new extender is ideal for your e-commerce and Web self-service applications. Net Search Extender searches text data contained in DB2 and can handle the heavy text search demands of large, text-intensive Web sites. It rapidly searches data without locking database tables and delivers excellent query performance and scalability.

Functions available in Net Search Extender

Net Search Extender delivers the following capabilities:

Indexing

At index creation time, you can use a command line interface to create your search indexes in main memory. Flexible index options permit the use of the same text column in one or more indexes. Net Search indexing proceeds without placing locks on the text data.

Searching

You invoke text searches via a stored procedure interface on your server. The powerful search algorithms permit searching by word, phrase, stem, or by fuzzy search. Tags define limited sections of the text document to restrict the range of the search. You can combine conditions for the search argument, using Boolean expression and wildcard operations.

Manageable search results

You specify sort definitions for the search results during index creation. Use result subsets and limits on search terms for your search to manage search performance where large data volumes might be involved. Cursor positioning permits you to navigate through the result set.

In-memory technology

At index creation time, you specify the parts of the table that you want to store in main memory, such as the table columns that you want to sort for a text index. For example, you could presort a text index on a book abstract column by the value in the price column.

Here is a hypothetical example of how Net Search Extender can perform these functions. A customer accesses an online bookstore on the Web that has information about their merchandise in a DB2 for OS/390 and z/OS database. The customer requests a list of all books about relational database. The DB2 application responds with a list that is ordered by price and includes the name of the author. Net Search Extender completes the transaction as shown in Figure 16 on page 43.

- The Web site customer enters their query into the order application
- The application does a stored procedure call to the database
- Net Search issues the text query "relational databases" against the text indexes, saving time and locks by not searching through standard SQL query

- The search engine preserves the order specified during index creation in the result list
- Net Search Extender returns corresponding results from the main memory table, listing the lowest priced books first

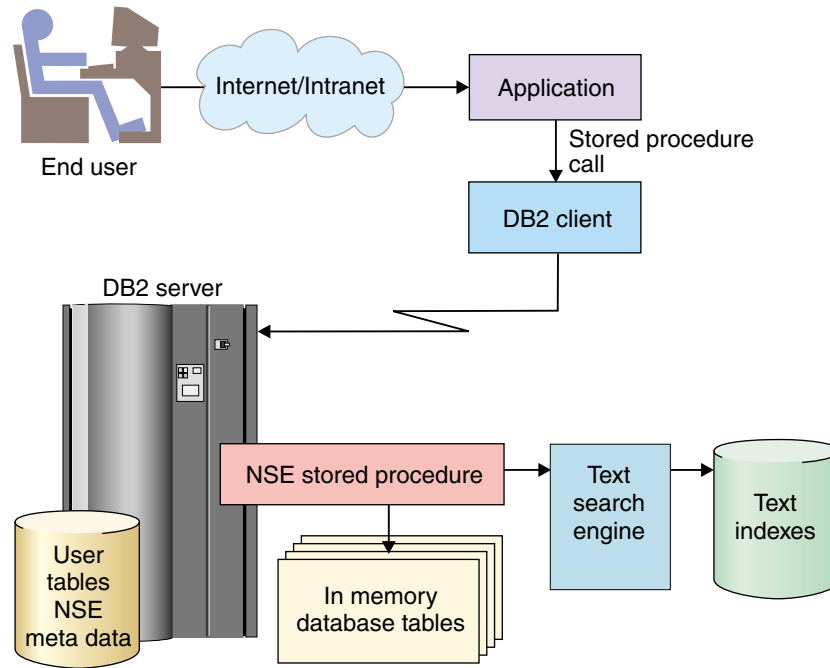


Figure 16. Searching a Web site using Net Search Extender

Using Net Search Extender

The powerful technology of Net Search Extender complements the search function in Text Extender. Net Search Extender provides a single stored procedure for searching large amounts of text. The stored procedure lets you use DB2 interfaces between your application and the DB2 server. However, you cannot combine your text search with a parametric search, as in a standard SQL query.

DB2 Text Extender provides text search function that is fully integrated into SQL. You can imbed these functions inside a standard SQL statement, such as inside the built-in SQL functions LENGTH and CONCAT. You can perform linguistic, proximity searches. For example, you can search for a sentence that includes both the word *big* and the word *small*. Use Text Extender with your XML documents, for thesaurus support, and to perform integrated multimedia searches, such as searching images.

Net Search is ideal for searching Web sites that are text-intensive and that receive a high level of concurrent queries. You can use the stored procedure in Net Search Extender with any Net.Data, Java, or DB2 CLI client application.

Net.Data for secure Web applications

Net.Data takes advantage of the System/390 capabilities as a premier platform for electronic commerce and Internet technology. Net.Data Web applications provide continuous application availability, scalability, security, and high performance. Using Net.Data for OS/390 for Web-based applications in the enterprise, you can:

Publish interactive data and provide universal access to dynamic data

Web applications need access to the key data in the enterprise. Net.Data provides dynamic access to DB2 family data, DRDA-enabled data sources, and HFS (hierarchical file system) flat-file data.

Benefit from high-performance Web-enabled data applications

Net.Data on OS/390 benefits from the proximity of Web applications to the DB2 data source. Net.Data eliminates the need to transfer large volumes of data over a network to a workstation server. Net.Data applications can call DB2 stored procedures to take advantage of DB2 performance and to reduce network traffic.

Rely on System/390 availability, scalability, and security

Net.Data relies on the strength of the System/390 platform as a highly available, scalable network server. Net.Data works with DB2 and IBM's HTTP Server for secure Web applications.

Use the robust application development function

Net.Data provides a powerful macro capability for robust Web application development. The Net.Data application can call DB2 stored procedures, as well as scripts and programs that are written in Perl, REXX, and C/C++. The application can also call Java applets and JavaScripts. Net.Data also provides an ODBC interface for DB2.

Net.Data Version 7 offers all of the capability of previous releases of Net.Data and much more, including many new features and performance, scalability, tracing, and service enhancements. Using Version 7, you can :

- Use Net.Data as a FastCGI application
- Use Support for Web page caching as well as the manual management of LOBs and cached Web pages when Net.Data is configured for CGI
- Upload files to the server
- Generate XML compliant documents using the new XMLBlock
- Call SQL functions from the REPORT and ROW blocks of other SQL functions
- Use a new language environment for running COBOL applications
- Use several new Net.Data built-in functions
- Write user-specified messages to the Net.Data error log and the Net.Data trace log through built-in functions and user-written Language Environments

Data Management Tools for your database

IBM is investing heavily in developing tools for your DB2 databases. Now you can select the tools that you need at a price that you like from a group of products that will help you meet your business challenges. A wide variety of tools are available:

IBM DB2 Administration for OS/390

A tool that provides a comprehensive set of database management functions and includes two new functions. ALTER supports modification of tables and their attributes. MIGRATE provides the facilities to copy data and objects to other DB2 subsystems.

Program number: 5655-E70

IBM DB2 Archive Log Compression Tool for z/OS

Restores data directly from highly compressed, off-site copies of DB2 logs for a complete disaster recovery support system. Database administrators can reduce the volume of archived logs, resulting in shorter I/O and recovery times. The tool also helps lower storage costs by helping you store logs on affordable direct access storage device (DASD) instead of tape.

Program number: 5655-F54

IBM DB2 Automation Tool for z/OS

Continuously and automatically coordinates the execution of DB2 tools, so you can realize the full potential of your DB2 system. Capabilities include the automatic execution of DB2 tools against specified objects; manual, periodic or rules-based execution of any number of tools; easy creation of job specifications; and development of job profiles through intuitive interactive system productivity facility (ISPF) panels.

Program number: 5697-G63

IBM DB2 Bind Manager

Determines if a bind is required after an application has been precompiled. If not, the tool automatically resets the time stamp and bypasses the bind. The DBRM Checker function handles consistency checking between an existing DBRMLIB and a DB2 subsystem. The Path Checker function lets you quickly determine whether a bind of a DBRM will result in a changed access path.

Program number: 5655-D38

IBM DB2 Change Accumulation Tool for z/OS

Restores database objects in the most precise and least disruptive manner for point-in-time recovery of database objects. The tool allows recovery routines to focus on single objects and previous states, and it produces SHRLEVEL REFERENCE image copies without the associated overhead and data locking.

Program number: 5655-F55

IBM DB2 DataPropagator for OS/390

Lets you maintain consistent copies of relational data across the DB2 family of databases, and automatically capture and apply data changes.

Program number: 5655-E60

IBM DB2 High Performance Unload for OS/390

Performs sequential reading and accessing of DB2 data at top speed. DB2 HP Unload can scan a table space and create output files in the format you need. It also lets you unload partitioned table spaces in parallel.

Program number: 5655-E69

IBM DB2 Log Analysis Tool for OS/390

Ensures high availability and complete control over data integrity. You can monitor and locate data changes, easily reverse undesired changes, and carry out re-do operations with precision.

Program number: 5655-E66

IBM DB2 Object Comparison Tool for z/OS

Compare objects and dependent objects from one source to those in another, so your test and development database system can be a mirror image of the production system. The tool runs as an extension to DB2 Administration Tool for OS/390, Version 2 Release1. The tool includes a

DB2 catalog extract function, a batch job generator, a dynamic link library (DDL) extract feature, an ISPF interface, and a batch compare capability.

Program number: 5697-G64

IBM DB2 Object Restore Tool for OS/390

Automatically restores previously dropped objects and all related dependencies. This tool saves DASD space because it eliminates the need for a duplicate shadow copy of the catalog to recover objects.

Program Number: 5655-E72

IBM DB2 Performance Monitor for OS/390

Helps you monitor, analyze and optimize your DB2 system performance. You can now take advantage of Data Sharing (Sysplex) Monitoring Online with group scope view, dynamic SQL statement cache monitoring, and an API that easily retrieves performance data and passes it on to an application program.

Program number: 5655-E61

IBM DB2 Query Monitor for OS/390

Helps you maximize DB2 availability. DB2 Query Monitor lets you determine what data is gathered during activity monitoring, when it is gathered, what database resources are required, and what kind of alerts or corrective actions to take.

Program number: 5655-E67

IBM DB2 Recovery Manager for OS/390

Simplifies and coordinates the recovery of both DB2 and IMS data to a common point, cutting the time and cost of data recovery and availability. DB2 Recovery Manager eliminates the error-prone complexity of managing different logs, utilities, and processes to do recovery from both databases.

Program Number: 5697-F56

IBM DB2 Row Archive Manager for OS/390

Saves storage, improves performance and reduces the overall costs of your DB2 environment by providing a simple method to control the separation of aged data from active DB2 data.

Program number: 5655-E65

IBM DB2 SQL Performance Analyzer for OS/390

Delivers performance analysis for all phases of database application design and development. DB2 SQL Performance Analyzer helps you find out how long queries will take and aids in preventing queries from running too long.

Program number: 5697-F57

IBM DB2 Table Editor for z/OS

Offers a multipurpose table-editing environment for enterprise-wide, direct update and data creation operations on DB2 Universal Database for z/OS from within Java, , Microsoft Windows or ISPF interfaces. This new generation of DB2 Forms for OS/390 enables database administrators and developers to conduct robust table editing in full-screen, multirow, single row, or wizard mode across multiple end-user environments and DB2 databases.

Program number: 5697-G65

IBM DB2 Utilities Suite for OS/390 and z/OS

Is comprised of DB2 Operational Utilities and DB2 Diagnostic and Recovery

Utilities. This suite extends the capabilities of the core DB2 for OS/390 and z/OS database utilities with advanced functions once available only through independent software vendors. The latest functional enhancements include object wildcarding and dynamic allocation, as well as a new UNLOAD utility and improvements to the LOAD and Online REORG utilities.

Program numbers:

DB2 Utilities Suite Program number: 5697-E98

DB2 Operational Utilities Program number: 5655-E63

DB2 Diagnostic and Recovery Utilities Program number: 5655-E62

IBM DB2 Web Query Tool for OS/390

Allows end users and administrators to access enterprise data with speed, reliability and simplicity. The tool enables pervasive connectivity over the Internet to every desktop.

Program number: 5655-E71

For more information about these tools and how to order them, see the Data Management Tools Web site:

ibm.com/software/data/db2imstools

Glossary

The following terms and abbreviations are defined as they are used in the DB2 library.

A

API. Application programming interface.

application plan. The control structure that is produced during the bind process. DB2 uses the application plan to process SQL statements that it encounters during statement execution.

ASCII. An encoding scheme that is used to represent strings in many environments, typically on PCs and workstations. Contrast with *EBCDIC* and *Unicode*.

B

buffer pool. Main storage that is reserved to satisfy the buffering requirements for one or more table spaces or indexes.

C

call level interface (CLI). A callable application programming interface (API) for database access, which is an alternative to using embedded SQL. In contrast to embedded SQL, DB2 ODBC (which is based on the CLI architecture) does not require the user to precompile or bind applications, but instead provides a standard set of functions to process SQL statements and related services at run time.

catalog. In DB2, a collection of tables that contains descriptions of objects such as tables, views, and indexes.

CCSID. Coded character set identifier.

character set. A defined set of characters.

character string. A sequence of bytes that represent bit data, single-byte characters, or a mixture of single-byte and multibyte characters.

clause. In SQL, a distinct part of a statement, such as a SELECT clause or a WHERE clause.

CLI. Call level interface.

client. See *requester*.

coded character set. A set of unambiguous rules that establish a character set and the one-to-one relationships between the characters of the set and their coded representations.

code page. A set of assignments of characters to code points. In EBCDIC, for example, the character 'A' is assigned code point X'C1', and character 'B' is assigned code point X'C2'. Within a code page, each code point has only one specific meaning.

code point. In CDRA, a unique bit pattern that represents a character in a code page.

collection. A group of packages that have the same qualifier.

column. The vertical component of a table. A column has a name and a particular data type (for example, character, decimal, or integer).

commit. The operation that ends a unit of work by releasing locks so that the database changes that are made by that unit of work can be perceived by other processes.

commit point. A point in time when data is considered consistent.

created temporary table. A table that holds temporary data and is defined with the SQL statement CREATE GLOBAL TEMPORARY TABLE. Information about created temporary tables is stored in the DB2 catalog, so this kind of table is persistent and can be shared across application processes. Contrast with *declared temporary table*. See also *temporary table*.

cursor. A named control structure that an application program uses to point to a row of interest within some set of rows, and to retrieve rows from the set, possibly making updates or deletions.

cursor sensitivity. The degree to which database updates are visible to the subsequent FETCH statements in a cursor. A cursor can be sensitive to changes that are made with positioned update and delete statements specifying the name of that cursor. A cursor can also be sensitive to changes that are made with searched update or delete statements, or with cursors other than this cursor. These changes can be made by this application process or by another application process.

D

DASD. Direct access storage device.

database. A collection of tables, or a collection of table spaces and index spaces.

database request module (DBRM). A data set member that is created by the DB2 precompiler and that contains information about SQL statements. DBRMs are used in the bind process.

data sharing • SQL function

data sharing. The ability of two or more DB2 subsystems to directly access and change a single set of data.

data sharing group. A collection of one or more DB2 subsystems that directly access and change the same data while maintaining data integrity.

data type. An attribute of columns, literals, host variables, special registers, and the results of functions and expressions.

DBCLOB. Double-byte character large object.

DBCS. Double-byte character set.

DB2 catalog. Tables that are maintained by DB2 and contain descriptions of DB2 objects, such as tables, views, and indexes.

declared temporary table. A table that holds temporary data and is defined with the SQL statement DECLARE GLOBAL TEMPORARY TABLE. Information about declared temporary tables is not stored in the DB2 catalog, so this kind of table is not persistent and can only be used by the application process that issued the DECLARE statement. Contrast with *created temporary table*. See also *temporary table*.

Distributed Relational Database Architecture (DRDA). A connection protocol for distributed relational database processing that is used by IBM's relational database products. DRDA includes protocols for communication between an application and a remote relational database management system, and for communication between relational database management systems.

double-byte character set (DBCS). A set of characters, which are used by national languages such as Japanese and Chinese, that have more symbols than can be represented by a single byte. Each character is 2 bytes in length. Contrast with *single-byte character set* and *multibyte character set*.

DRDA. Distributed Relational Database Architecture.

G

GB. Gigabyte (1 073 741 824 bytes).

K

KB. Kilobyte (1024 bytes).

L

list. A type of object, which DB2 utilities can process, that identifies multiple table spaces, multiple index spaces, or both. A list is defined with the LISTDEF utility control statement.

log. A collection of records that describe the events that occur during DB2 execution and that indicate their sequence. The information thus recorded is used for recovery in the event of a failure during DB2 execution.

M

MB. Megabyte (1 048 576 bytes).

multibyte character set (MBCS). A character set that represents single characters with more than a single byte. Contrast with *single-byte character set* and *double-byte character set*. See also *Unicode*.

O

ODBC. Open Database Connectivity.

Open Database Connectivity (ODBC). A Microsoft database application programming interface (API) for C that allows access to database management systems by using callable SQL. ODBC does not require the use of an SQL preprocessor. In addition, ODBC provides an architecture that lets users add modules called *database drivers*, which link the application to their choice of database management systems at run time. This means that applications no longer need to be directly linked to the modules of all the database management systems that are supported.

S

savepoint. A named entity that represents the state of data and schemas at a particular point in time within a unit of work. SQL statements exist to set a savepoint, release a savepoint, and restore data and schemas to the state that the savepoint represents. The restoration of data and schemas to a savepoint is usually referred to as *rolling back to a savepoint*.

scrollable cursor. A cursor that can be moved in both a forward and a backward direction.

self-referencing constraint. A referential constraint that defines a relationship in which a table is a dependent of itself.

self-referencing table. A table with a self-referencing constraint.

single-byte character set (SBCS). A set of characters in which each character is represented by a single byte. Contrast with *double-byte character set* or *multibyte character set*.

SQL function. A user-defined function in which the CREATE FUNCTION statement contains the source code. The source code is a single SQL expression that evaluates to a single value. The SQL user-defined function can return only one parameter.

SQLJ. Structured Query Language (SQL) that is embedded in the Java® programming language.

star join. A method of joining a dimension column of a fact table to the key column of the corresponding dimension table. See also *join*, *dimension*, and *star schema*.

star schema. The combination of a fact table (which contains most of the data) and a number of dimension tables. See also *star join*, *dimension*, and *dimension table*.

stored procedure. A user-written application program that can be invoked through the use of the SQL CALL statement.

Structured Query Language (SQL). A standardized language for defining and manipulating data in a relational database.

T

table. A named data object consisting of a specific number of columns and some number of unordered rows. See also *base table* or *temporary table*.

TB. Terabyte (1 099 511 627 776 bytes).

template. A DB2 utilities output data set descriptor that is used for dynamic allocation. A template is defined by the TEMPLATE utility control statement.

temporary table. A table that holds temporary data; for example, temporary tables are useful for holding or sorting intermediate results from queries that contain a large number of rows. The two kinds of temporary table, which are created by different SQL statements, are the created temporary table and the declared temporary table. Contrast with *result table*. See also *created temporary table* and *declared temporary table*.

type 2 indexes. Indexes that are created on a release of DB2 after Version 6 or that are specified as type 2 indexes in Version 4 or later.

U

Unicode. A standard that parallels the ISO-10646 standard. Several implementations of the Unicode standard exist, all of which have the ability to represent a large percentage of the characters contained in the many scripts that are used throughout the world.

unit of recovery. A recoverable sequence of operations within a single resource manager, such as an instance of DB2. Contrast with *unit of work*.

unit of work. A recoverable sequence of operations within an application process. At any time, an application process is a single unit of work, but the life of an application process can involve many units of

work as a result of commit or rollback operations. In a *multisite update* operation, a single unit of work can include several *units of recovery*. Contrast with *unit of recovery*.

Z

z/OS. An operating system for the eServer product line that supports 64-bit real storage.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J74/G4
555 Bailey Avenue

P.O. Box 49023
San Jose, CA 95161-9023
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both.

AIX	IBM
AS/400	IMS
BookManager	IMS/ESA
BookMaster	MQSeries
CICS	Net.Data
CICS/ESA	OS/2
DATABASE 2	OS/390
DataJoiner	OS/400
DataPropagator	Parallel Sysplex
DB2	QMF
DB2 Connect	RACF
DB2 OLAP Server	RAMAC
DB2 Universal Database	S/390
DRDA	System/390
Enterprise Storage Server	WebSphere

Tivoli is a trademark of Tivoli Systems Inc. in the United States, other countries, or both.

Java Beans, JDBC, and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and Visual Basic are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Index

A

AIX 3
application enhancements 16
applications
 Web 42
 Net Search Extender 42
Automatic Restart Manager 14
availability enhancements 10

B

Baan 23
books, new 25
business intelligence 1
business partnerships 23

C

cancel thread without rollback processing 13
CASTOUT(NO) option on -STOP DB2 command 15
CHECKPAGE 14
collections with the DB2 UDB Control Center 36
commands 13
commands with a global scope 15
COMMIT and ROLLBACK in stored procedures 23
CONNECT statement 19
constraints 10
COPY, external 14
COPY with CHECKPAGE 14
COPYTOCOPY utility 11
customer relationship management 1

D

data definitions, generating using DB2 UDB Control Center 35
Data Management Tools 44
data set templates 36
data sharing enhancements 14
DB2 Administration Tool 35
DB2 Connect 3
DB2 Estimator 41
DB2 Extenders 22
DB2 Installer 38
DB2 Management Clients package 34
DB2 Stored Procedures Builder 37
DB2 UDB Control Center 34
DB2 Visual Explain 39
DB2 Warehouse Manager 5, 29
DBADM authority 10
DDF RESUME command 14
DDF SUSPEND command 14
declared temporary tables 21
deferred definitions of data sets 10
defining identity columns 21
DELETE or UPDATE with self-referencing subselect 18
distributed applications
 coordinating updates 4

distributed applications (*continued*)
 global transaction support 20
 monitoring 20
 performance improvement 19
dynamic allocation of data sets 9

E

e-business 1
 Net.Data 44
 Net Search Extender 42
 scalability 10
 system parameters, changing online 13
 Unicode support 22
 Version 7 enhancements 5, 9
 Web sites 26
 XML extender 16
enterprise resource management 1, 23
Enterprise Storage Server (ESS) 16
extenders 22

F

family compatibility 5, 9, 18
fast implicit close 19
features, DB2 for OS/390 and z/OS 29
FETCH FIRST *n* ROWS SQL clause 19

G

global transaction support 4
global transaction support for distributed applications 20

I

identity columns 21
IFI calls with global scope 15
independent software developers 23
index access for small tables 14
Information Center 37

J

Java
 Control Center run as a Java application 34
 Java applets and Javascript with Net.Data 44
 Java stored procedures 20
 reports with Java query capability 31
 stored procedures 37

K

Kerberos security 23

L

LISTDEF facility 9
LOAD RESUME, online 13

LOAD utility, enhancements 20
log read request, retry 14
LOG SUSPEND and RESUME commands 14

M

monitoring
 distributed applications 20
 server elapsed time for remote requests 24
multiple inputs to parallel LOAD 11
multisite updates 4

N

Net.Data 3, 44
Net Search Extender 5, 42
notices, legal 53

O

object lists 36
ODBC 21
optional features, DB2 for OS/390 and z/OS 29
OS/390 Workload Manager (WLM) 3
OS/400 3

P

parallel LOAD with multiple inputs 11
pattern-matching characters in utility lists 9
PeopleSoft 10, 23
performance
 Enterprise Storage Server and DB2 16
 estimating 41
Persistent Stored Modules 21
persistent structure size 15
precompiler servers 22
publications, new 25

Q

QMF 31
QMF for Windows 32
QMF High Performance Option (HPO) 31
QMF HPO 31
queries using a new join method 13
query parallelism 12

R

remote request response times 24
REORG, online 13
restart light for data sharing members 14
REXX 21

S

SAP 10, 23
savepoints 21
scalability improvements 10
scalar function 19

58 What's New?

scrollable cursors 18
security
 Web with Net.Data 44
 Windows Kerberos 23
 workstation clients 23
Seibel Systems 23
self-referencing subselect on UPDATE and DELETE 18
server elapsed time monitoring 24
Shutdown of data sharing members 15
Solaris Operating Environment 3
SQL
 CONNECT statement 19
 embedded, support for scrollable cursors using 18
 enhancements 18
 processing queries 12
 scalar function 19
 SELECT statement, ORDER BY 19
 support for COMMIT and ROLLBACK in stored procedures 23
 support for REXX 21
 with XML extenders 17
SQL procedures 21
star schema 13
statistics history, comprehensive 9
stored procedures
 DB2 Stored Procedures Builder 37
 Java stored procedures 20
 support for COMMIT and ROLLBACK 23
 written entirely in SQL 21
structure size changes, persistent 15
system parameters, changing online 13

T

tables
 declared temporary, defining 21
 small, with index access 14
 UNION and UNION ALL operators, using 18
templates, data set 36
Text Extender 22
three-tier architecture 1
transaction manager support in DB2 4

U

Unicode support 22
UNION and UNION ALL operators 18
UNLOAD utility 10
UPDATE or DELETE with self-referencing subselect 18
UPDATE with subselect 19
utilities
 COPY 14
 COPYTOCOPY 11
 DB2 UDB Control Center, using with 35
 dynamic allocation of data sets 9
 LOAD 11
 LOAD RESUME, online 13
 procedures using the DB2 UDB Control Center 36
 REORG 13
 UNLOAD 10

utilities (*continued*)
 using pattern-matching characters 9

W

warehouse management 29

Web

- accessing books 37
- applications 42
 - Net Search Extender 42
- publishing reports using QMF 31
- sites 26

WebSphere 3

Windows 2000 3

Windows Kerberos security 23

Windows NT 3

X

XML extender

- description 16

XML extenders

- using SQL 17



Program Number: 5675-DB2



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC26-9946-01

